

Principles of Appearance Acquisition and Representation

By Tim Weyrich, Jason Lawrence,
Hendrik P. A. Lensch,
Szymon Rusinkiewicz and Todd Zickler

Contents

1 Radiometry and Appearance Models	76
1.1 Radiometry	78
1.2 Surface Reflectance	82
1.3 6D Datasets: SVBRDFs, BTFs, and Distant-Light Reflectance Fields	91
1.4 Subsurface Scattering	92
1.5 8D Reflectance Fields	94
1.6 Generalizing Reflectance and Scattering	95
2 Principles of Acquisition	97
2.1 Traditional BRDF Measurement: The Gonioreflectometer	98
2.2 Image-based Acquisition	101
2.3 Acquiring BRDF, SVBRDF and 6D Reflectance Fields	109
2.4 Acquiring BSSRDF	116
2.5 Separating Reflection Components	120

3 Spatially- (and Temporally-) Varying Reflectance Models	129
3.1 Acquisition	130
3.2 Representation	132
3.3 Case Study: Basis Decomposition and the Inverse Shade Tree Framework	139
3.4 Space- and Time-Varying BRDFs (STVBRDFs)	144
4 From BSSRDF to 8D Reflectance Fields	146
4.1 Capturing BTFs	146
4.2 Acquiring BSSRDFs	147
4.3 Diffuse Subsurface Scattering	149
4.4 Arbitrary Light Transport	150
4.5 Time-Resolved Capture and Interpolation	157
4.6 Conclusions and Future Work	158
5 Acquiring Human Face Appearance	159
5.1 Skin Appearance	159
5.2 Face Acquisition Systems	162
5.3 Appearance Editing	164
5.4 Analysis of Human Face Appearance	165
6 Open Problems and Future Work	174
References	177

Principles of Appearance Acquisition and Representation

Tim Weyrich^{1,2}, Jason Lawrence³,
Hendrik P. A. Lensch⁴,
Szymon Rusinkiewicz⁵ and Todd Zickler⁶

¹ Princeton University, 35 Olden St, Princeton, NJ, 08540, USA,
tweyrich@cs.princeton.edu

² University College London, Gower St, London, WC1E 6BT, UK,
t.weyrich@cs.ucl.ac.uk

³ University of Virginia, 151 Engineer's Way, P.O. Box 400740,
Charlottesville, VA, 22904-4740, USA, jdl@cs.virginia.edu

⁴ Ulm University, Institute for Media Informatics, James-Franck-Ring,
89081 Ulm, Germany, hendrik.lensch@uni-ulm.de

⁵ Princeton University, 35 Olden St, Princeton, NJ, 08540, USA,
smr@cs.princeton.edu

⁶ Harvard University, 33 Oxford St, Cambridge, MA, USA,
zickler@eecs.harvard.edu

Abstract

Algorithms for scene understanding and realistic image synthesis require accurate models of the way real-world materials scatter light. This study describes recent work in the graphics community to measure the spatially- and directionally-varying reflectance and subsurface scattering of complex materials, and to develop efficient representations and analysis tools for these datasets. We describe the design of acquisition devices and capture strategies for reflectance functions such as BRDFs and BSSRDFs, efficient factored representations, and a case study of capturing the appearance of human faces.

1

Radiometry and Appearance Models

Comprehending the visual world around us requires understanding the role of *materials*. In essence, we think of the appearance of a material as being a function of how that material interacts with light. The material may reflect light or may exhibit more complex phenomena such as subsurface scattering.

Reflectance is itself a complex phenomenon. In general, a surface may reflect a different amount of light at each position, and for each possible direction of incident and exitant light (Figure 1.1, left). So, to completely characterize a surface's reflection we need a six-dimensional function giving the amount of light reflected for each combination of these variables (position and incident and exitant directions are two dimensions each). Note that this does not even consider such effects as time or wavelength dependence. We will consider those later, but for now let us simply ignore all time dependence and assume that any wavelength dependence is aggregated into three color channels: red, green, and blue.

These reflectance functions embody a significant amount of information. They can tell us whether a surface is shiny or matte, metallic or dielectric, smooth or rough. Knowing the reflectance function for

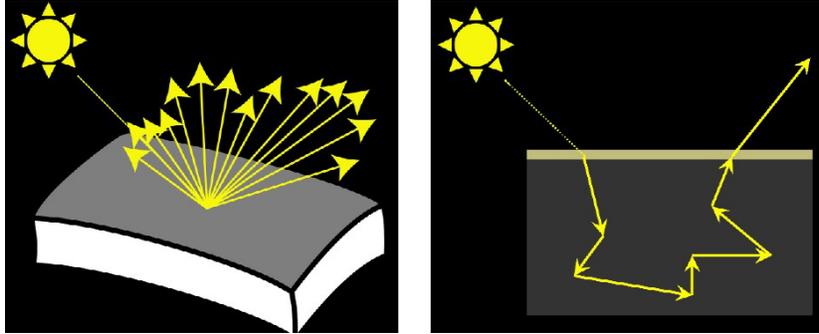


Fig. 1.1 Materials can exhibit reflectance (left), subsurface scattering (right), or more complex scattering phenomena.

a surface allows us to make complete predictions of how that surface appears under any possible lighting.

For translucent surfaces, the interaction with light can no longer be described as simple reflection. This is because light leaves the surface at a different point than where it entered (Figure 1.1, right). So, in order to characterize such surfaces we need a function that gives the amount of light that is scattered from each possible position (2D) to each other position (another 2D). To be even more correct, of course, it would be necessary to account for the directional dependency as well.

This study covers the basic principles of how materials are described, how the appearance of real-world objects may be measured, and how a knowledge of appearance aids in a variety of applications. In addition to the obvious application domain of image synthesis, having a complete knowledge of a material’s appearance can help in interpreting images. It will aid in 3D reconstruction, view interpolation, and object recognition. Furthermore, knowing how to characterize materials can help in understanding how humans perceive surfaces.

This section covers foundational topics. It will survey the domain of *radiometry* and introduce the definition of the Bidirectional Reflectance-Distribution Function (BRDF): a function describing surface reflectance at a point. It will then cover generalizations of the BRDF, including spatial variation and subsurface scattering. Finally, it will consider the many different types of data that can be captured that characterizes “appearance,” and how they relate to each other.

1.1 Radiometry

The field of radiometry is concerned with the characterization of the “amount” of electromagnetic radiation, including light, flowing in space. Though this chapter presents some fundamental concepts, the reader is referred to classic works such as those of Ishimaru [80] for more details.

To begin, it is necessary to consider the different quantities related to light flow, and the *radiometric units* in which they are expressed. Light is a form of electromagnetic energy, and so can be measured using the SI units of Joules. Because in graphics and vision we usually consider steady-state flows, instead of individual pulses or quanta, we will most often be interested in the amount of energy flowing per unit time. This is known as “radiant flux” (Φ) or just “power,” and hence may be measured using the SI units of Watts.

Although having a way of characterizing the total flow of light is useful, we will need to consider more complex quantities in order to talk about concepts such as light sources and surface reflectance.

Point Light in a Direction: Consider an ideal light source (idealized as a point in space). If the light were being emitted uniformly in all directions, describing its power (in Watts) would characterize it completely. However, it is possible that light is not being emitted equally in all directions. In this case, characterizing the power being emitted in a *particular* direction requires a different unit. In such cases, we can talk about the amount of power being emitted per unit *solid angle*.

So what exactly is a solid angle, and how is it measured? A useful analogy is to the way an angle is defined in the plane. One radian is defined as the angle subtended by an arc of a circle, with the arc length being equal to the circle’s radius. Equivalently, an angle in radians may be calculated by dividing the length of a circular arc by the radius.

Moving to the concept of solid angles, we will be working in three dimensions (vs. two for angles), and will be looking at a sphere (vs. a circle). The fundamental unit of solid angle is known as the *steradian*, and is defined as the area of some region on a sphere divided by the

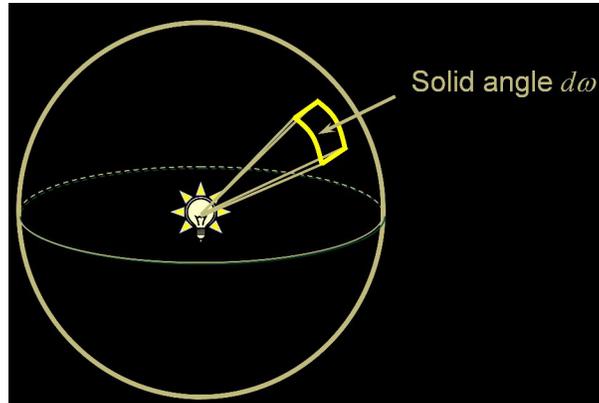


Fig. 1.2 Point light source emitting light in a direction.

square of the sphere’s radius. A complete sphere thus has 4π steradians, and smaller solid angles define smaller regions of the space of directions.

So, measuring the directional power or *radiant intensity* of a point light source can be done using the units of Watts per steradian:

$$I = \frac{d\Phi}{d\omega}. \quad [\text{W} \cdot \text{sr}^{-1}] \quad (1.1)$$

The same amount of power emitted into a smaller solid angle will result in a larger measurement (e.g., consider a laser, which has relatively low power but concentrated into a small solid angle).

Light Falling on a Surface: Another radiometric quantity we often wish to measure is called *irradiance*. It represents the amount of light falling onto a surface. Because the same radiant flux will be “more concentrated” when falling onto a smaller area of surface than a larger surface, we define irradiance E as power per unit area:

$$E = \frac{d\Phi}{dA}. \quad [\text{W} \cdot \text{m}^{-2}] \quad (1.2)$$

Note that we write this definition in differential form, to emphasize that we are concerned with the limit of incident power per unit area, as that area shrinks to zero.

Given this definition of irradiance, there are two immediate and easily-observed “laws” that emerge. The first is the inverse-square

law: moving a point light source away from a surface reduces irradiance in proportion to the inverse square of the distance. Secondly, tilting a surface away from a point light results in a lower irradiance, in proportion to the cosine of the angle between the surface normal and the direction toward the light. This “cosine law” is often written as the dot product between the (unit-length) surface normal and light vectors.

Light Emitted from a Surface in a Direction: We now come to the final, and most complex, radiometric quantity we are going to consider, which describes the emission of light from a surface. This can be thought of as combining the two concepts we just saw: the emitted light can vary with direction (hence we must control for its directional distribution, as we did with the point-light case), and we are interested in the amount of light emitted per unit surface area. This is almost enough for a practical definition of *radiance*, but it is conventional to use a slightly different, “observer-based” definition of surface area, instead of the one used for irradiance. In particular, an observer or sensor measuring light emitted from a surface will be sensitive to *projected* surface area, perpendicular to the viewing direction (see Figure 1.3).

Hence, we arrive at the definition of radiance: power emitted per unit projected area (perpendicular to the viewing direction) per unit

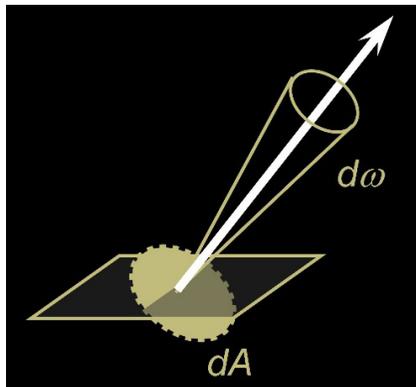


Fig. 1.3 Radiance is defined as light emitted from a surface, in a specific direction, per unit (projected) area.

solid angle:

$$L = \frac{d\Phi}{dA_{\text{proj}} d\omega}. \quad [\text{W} \cdot \text{m}^{-2} \cdot \text{sr}^{-1}] \quad (1.3)$$

Radiance is perhaps the most fundamental unit in computer vision and graphics. It is easy to show that the irradiance on a camera sensor is proportional to the radiance of the surfaces it is imaging, with the constant of proportionality determined by the imaging optics. (More accurately, the optical system effectively integrates the radiance over the solid angle subtended by the aperture, as seen from the surface.) The sensor irradiance at each pixel is converted to an electrical signal, then digitized, and so the pixel values we deal with in digital images are (ignoring effects such as gamma applied to the pixel values) proportional to radiance.

Integrating radiance over all exitant angles, including a cosine term to account for projected area, gives a quantity called *radiant exitance*, which is frequently encountered in graphics simulations:

$$M = \int_{\Omega} L(\theta, \varphi) \cos\theta d\omega. \quad [\text{W} \cdot \text{m}^{-2}] \quad (1.4)$$

When radiance is equal for all exitant directions, as is the case for some surfaces, this quantity is usually called *radiosity* and is conventionally denoted by the symbol B .

The Plenoptic Function and the Light Field: Radiance in a scene may be represented by the *plenoptic function*, which is a positive function defined on a five-dimensional domain:

$$L(x, y, z, \theta, \phi) \quad [\text{W} \cdot \text{m}^{-2} \cdot \text{sr}^{-1}]. \quad (1.5)$$

representing the radiance in every ray direction at every point in three-dimensional space [1]. Since radiance is constant along rays in free space, we can often drop one of these dimensions, and the resulting four-dimensional entity is referred to as the *lumigraph* [60] or, more commonly, the *(4D) light field* [105]. (Note that the term “light field” was originally introduced by Gershun [53] to describe a vector, rather than scalar, version of the 5D function.) The set of rays representing

a light field may be parameterized in several ways, in addition to the obvious point/angle parameterization $L(x, y, \theta, \phi)$. In particular, it is common to parameterize a “light slab” by the positions of ray intersections with two planes: $L = L(u, v, s, t)$, where (u, v) and (s, t) are the coordinates on two specified planes.

Radiometry vs. Photometry: The preceding discussion has focused purely on physical (radiometric) units, which is appropriate when dealing with acquisition apparatus. However, there is a parallel set of *photometric* units, which also take into account the intensity perceived by a human observer. In particular, they account for the fact that the human eye is sensitive to a range of wavelengths from 400 (blue) to 700 (red) nanometers, but that the sensitivity is not constant within that range.

The original photometric unit was an “international standard candle,” defined in terms of carbon filament lamps. Today, the *candela* is one of the seven base SI units: one candela is the luminous intensity of a light source producing 1/683 Watt per steradian, at a frequency of 540×10^{12} Hz (corresponding to green light with a wavelength of approximately 555 nm). Beginning with this unit, it is possible to define concepts analogous to radiant flux, irradiance, and radiance, namely luminous power (measured in lumens, where one lumen is equal to one candela times one steradian), illuminance (measured in lux = lumens per square meter), and luminance (measured in nits = candelas per square meter or lux per steradian).

While it is important to be aware of the difference between radiometry and photometry, we will assume single-wavelength, radiometric measurements in the remainder of this section.

1.2 Surface Reflectance

Having learned about radiometry, we are now ready to define the Bidirectional Reflectance-Distribution Function (BRDF), which characterizes reflection at a point on a surface [142]. Formally, it is the ratio between the reflected radiance of a surface and the irradiance that caused that reflection. The radiance and irradiance are each measured

at a particular angle of exitant and incident light, respectively, so the BRDF is usually written as a function of four variables: the polar angles of incident and exitant light.

$$f_r(\boldsymbol{\omega}_i \rightarrow \boldsymbol{\omega}_o) = f_r(\theta_i, \varphi_i, \theta_o, \varphi_o) = \frac{dL_o(\boldsymbol{\omega}_o)}{dE_i(\boldsymbol{\omega}_i)}. \quad [\text{sr}^{-1}] \quad (1.6)$$

The BRDF has units of inverse steradians and is often written as a differential quantity. This is to emphasize that there is no such thing as light arriving from *exactly* one direction, and being reflected into exactly one outgoing direction. Rather, we must look at non-zero incident and exitant solid angles, and consider the limit as those approach zero.

Because BRDFs are 4D functions, they are a bit tricky to visualize directly. Instead, we often visualize two-dimensional slices of this function. Figure 1.4 shows two 2D slices of a BRDF, each corresponding to one direction of incidence (the arrow) and all possible directions of reflection. The blue surface is a hemisphere stretched such that its radius in any direction is the reflected radiance in that direction, and is known as a goniometric plot.

You will note that, for this particular BRDF, some of the incident light is reflected equally in all directions. This is the constant-radius (spherical) portion of the surface you see. However, there is also a bump in the surface, indicating that there is a concentrated reflection in one particular direction.

If we change the direction of incidence, we see that the constant portion of the function remains unchanged, but the position of the bump moves. In fact, the bump always appears near the direction of

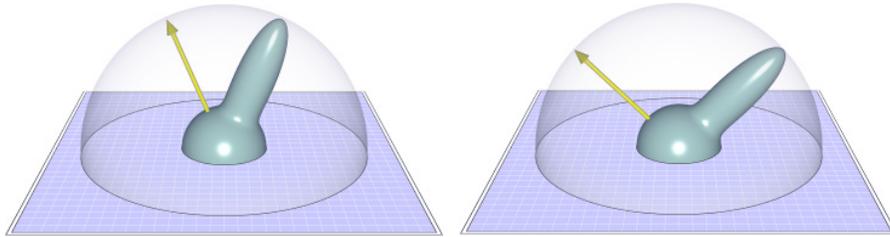


Fig. 1.4 Goniometric view of slices of a BRDF corresponding to two incident directions (denoted by the yellow arrows).

“ideal mirror reflection” of the incident direction. This is known as a specular highlight, and it gives a surface a shiny appearance.

Properties of the BRDF: Before we look at specific BRDF models, let us discuss a few properties shared by all BRDFs. The first is **energy conservation**. Because all incident light must be either reflected or absorbed, and no light may be created during reflection, it is impossible for a surface to reflect more light than was incident on it. Expressing this mathematically, we see that the integral of the BRDF over all outgoing directions, scaled by a cosine term to account for foreshortening, must be less than one:

$$\forall \omega_i : \int_{\Omega} f_r(\omega_i, \omega_o) \cos \theta_o d\omega_o \leq 1. \quad (1.7)$$

A second, more subtle, property of BRDFs is that they must be unchanged when the angles of incidence and exitance are swapped:

$$f_r(\omega_i \rightarrow \omega_o) = f_r(\omega_o \rightarrow \omega_i). \quad (1.8)$$

This is a condition known as **Helmholtz reciprocity**, and is due to the symmetry of light transport [186]. Some systems, such as the work on Helmholtz stereopsis [211], have relied on this property, which is often expressed as camera/projector duality: in many imaging systems it is possible to interchange the roles of camera and projector, provided that cosine terms are properly accounted for.

Though all real BRDFs satisfy the above two properties, measured data (which can include non-local effects) and the adhoc shading models used in graphics and vision frequently do not. The term *physically-plausible BRDF* is sometimes used for reflectance functions that satisfy energy conservation and reciprocity.

Some, but not all, BRDFs have a property called **isotropy**: they are unchanged if the incoming and outgoing vectors are rotated by the same amount about the surface normal. With isotropy, a useful simplification may be made: the BRDF is really a three-dimensional function in this case, and depends only on the *difference* between the azimuthal angles of incidence and exitance.

The inverse of isotropy is **anisotropy**. An anisotropic BRDF does not remain constant when the incoming and outgoing angles are

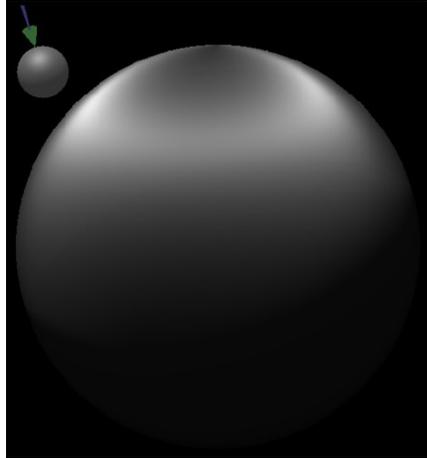


Fig. 1.5 Anisotropic reflection.

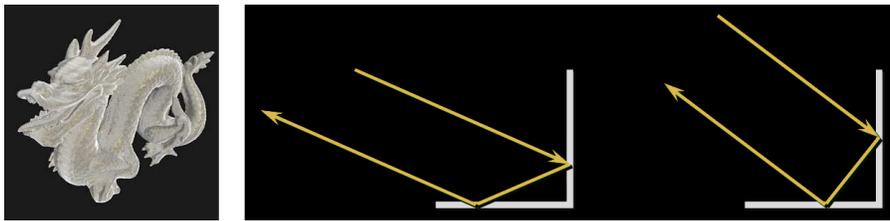


Fig. 1.6 Left: dusty surfaces exhibit an increase in reflection toward grazing angles [152]. Right: corner-reflectors are one example of a configuration that may produce retroreflection.

rotated. In this case, a full four-dimensional function is necessary to characterize the behavior of the surface. Anisotropic materials are frequently encountered when the surface has a strongly directional structure at the small scale: brushed metals are one example (Figure 1.5).

Another commonly observed characteristic of some BRDFs is **asperity scattering**: an increase in light reflected into all grazing angles, as is typical for “dusty” surfaces (Figure 1.6, left). Finally, some BRDFs exhibit **retro-reflection**. That is, they scatter light most strongly back into the direction from which it arrived. Street signs and the paint found on roads are common examples of this phenomenon, which is created through “corner reflector” configurations (Figure 1.6, right) or particles of high-index material embedded in paint.

Parameterization: Thus far, we have assumed that the 4D BRDF domain is parameterized by the spherical coordinates of the incident and reflected directions. We are free to choose any parameterization, of course, and there are others with significant advantages. We may require a parameterization without singularities, for example, or we may want one that allows a more compact or intuitive representation.

One useful parameterization of the BRDF [159] uses the “halfway” vector \mathbf{h} (i.e., the vector halfway between the incoming and reflected rays) and a “difference” vector \mathbf{d} , which is just the incident ray in a frame of reference in which the halfway vector is at the north pole (see Figure 1.7). Using the spherical angles of \mathbf{h} and \mathbf{d} , a point in the BRDF domain is written:

$$(\theta_h, \varphi_h, \theta_d, \varphi_d) \subset [0, 2\pi) \times [0, \pi/2) \times [0, \pi) \times [0, \pi/2). \quad (1.9)$$

A typical BRDF varies slowly over much of its domain, and the halfway/difference parameterization exploits this by moving the coordinate axes away from these regions. The axes are aligned with directions of common BRDF phenomena (specular and retro-reflective peaks) and this enables representations that are both intuitive and efficient.

Isotropy and Helmholtz reciprocity are conveniently described using the halfway/difference parameterization. Helmholtz reciprocity implies that the BRDF is unchanged under $\varphi_d \rightarrow \varphi_d + \pi$, so that φ_d can be restricted to $[0, \pi)$. Isotropy implies that the BRDF is a constant

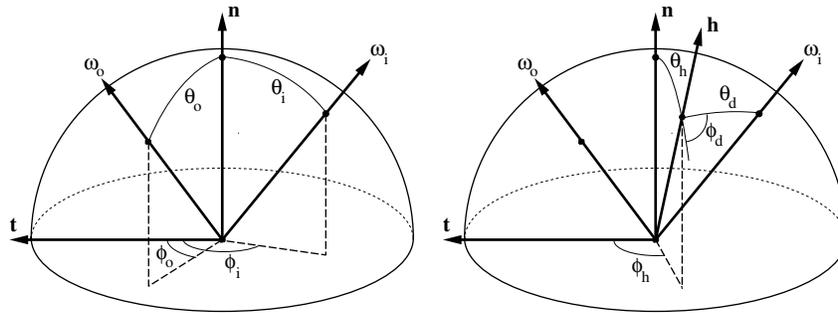


Fig. 1.7 Halfway/difference angle parameterization of BRDFs. Instead of treating the BRDF as a function of (θ_i, φ_i) and (θ_o, φ_o) , as shown at left, the BRDF is considered a function of the half-angle (θ_h, φ_h) and a difference angle (θ_d, φ_d) , as shown at right. The vectors marked \mathbf{n} and \mathbf{t} are the surface normal and tangent, respectively.

function of φ_h , meaning that this dimension can simply be ignored in the isotropic case.

For glossy surfaces, specular peaks occur at small half-angles (i.e., $\theta_h \approx 0$), but variation with respect to the difference angle (θ_d) is governed primarily by Fresnel reflection and tends to be limited for small and moderate values of θ_d .

Lambertian BRDF: We now turn to looking at specific examples of BRDFs. We will look at simple examples, such that the reflectance may be written as a mathematical formula. Real surfaces, of course, are more complex than this, and mathematical models frequently do not predict the reflectance with great accuracy.

The simplest possible BRDF is just a constant:

$$f_r = \text{const.} = \rho/\pi. \quad (1.10)$$

(Keep in mind that the BRDF is defined in terms of irradiance, which has the “incident cosine law” implicitly included.) This results in a matte or diffuse appearance, and is known as ideal Lambertian reflectance. This BRDF is frequently written as a constant ρ divided by π . In this case, ρ is interpreted as the diffuse albedo: it is the fraction of light that is reflected (vs. absorbed) by the surface. Plugging this BRDF into the energy conservation integral verifies that the surface conserves energy precisely when the albedo is less than or equal to one.

Phong and Blinn–Phong BRDFs: Another simple analytic BRDF is the Phong model [153], designed to qualitatively mimic the appearance of glossy materials:

$$f_r = k_s (\mathbf{r} \cdot \mathbf{v})^n, \quad (1.11)$$

where \mathbf{v} is the view direction and \mathbf{r} is the mirror reflection of the light direction from the tangent plane. Note that the Phong “BRDF” used in computer graphics often includes an additional $1/\cos\theta_o$ factor, which is canceled by the irradiance “cosine law.” This is not a physically-plausible BRDF: it does not exhibit reciprocity, and does not conserve energy.

A common variant of this model is sometimes known as the Blinn-Phong model [9]:

$$f_r = k_s (\mathbf{n} \cdot \mathbf{h})^n, \quad (1.12)$$

though again it is often stated as a physically-implausible shading model rather than a BRDF. Lewis [106] introduced a physically-plausible BRDF based on this model that is appropriately scaled to conserve energy.

In contrast to the Lambertian BRDF, the distribution of reflected light in these models is not constant. In fact, there is a lobe centered around the direction of ideal mirror reflection for each incident angle, containing significantly more energy than the rest of the domain. This is known as the specular lobe, and its size and width (fall-off) are controlled by the parameters k_s and n , respectively.

There are a few things to remember when working with the above models. First, they are not physically-based and only qualitatively reproduce the rough appearance of a specular lobe. Second, in computer graphics these models are frequently not presented as BRDFs, but rather operate on incident illumination quantities that have not had the “cosine law” applied. In this case, the models that are actually used are equivalent to “BRDFs” with the incident cosine divided out, and hence do not satisfy Helmholtz reciprocity. Finally, the specular exponents n in the original Phong and Blinn-Phong formulations are not equivalent in the widths of highlights they produce. To obtain roughly-equivalent highlights from the Blinn-Phong model, it is necessary to use an n that is four times as large as in the Phong model.

Lafortune BRDF: A popular model used for fitting analytic functions to measured BRDF data is the Lafortune model [93]:

$$f_r = (C_x l_x v_x + C_y l_y v_y + C_z l_z v_z)^n, \quad (1.13)$$

in which l_x , v_x , etc. represent the components of the light vector \mathbf{l} and view vector \mathbf{v} , in a coordinate system in which the surface normal is oriented along the z axis. This model reduces to Phong by choosing $-C_x = -C_y = C_z = \sqrt[n]{k_s}$, but through suitable choice of

parameters can also represent non-Lambertian diffuse reflection, off-specular reflection, anisotropy, and retro-reflection. It is also common to fit a sum of multiple lobes of (1.13) to measured datasets.

Ward BRDF: Another popular BRDF used in fits to measurements is the Ward model [191]:

$$f_r = k_s \frac{e^{-\tan^2 \theta_h ((\cos^2 \phi_h)/\alpha_x^2 + (\sin^2 \phi_h)/\alpha_y^2)}}{4\pi \alpha_x \alpha_y \sqrt{\cos \theta_i \cos \theta_o}}. \quad (1.14)$$

Compared to the Blinn–Phong BRDF, the Ward model includes a specular peak shaped by a Gaussian function (as opposed to a power-of-cosine model), but also can model anisotropic reflection by using separate Gaussian widths α_x and α_y in two perpendicular directions.

Torrance-Sparrow BRDF: Numerous BRDFs have been derived from first principles that predict the aggregate reflectance for surfaces that at a small scale consists of tiny, mirror-reflective “microfacets” oriented in random directions. An early microfacet BRDF was originally developed in the physics community by Torrance and Sparrow [181], introduced to the graphics community by Blinn [9], and later refined by Cook and Torrance [18]:

$$f_r = \frac{D G F}{\pi \cos \theta_i \cos \theta_o}. \quad (1.15)$$

There are three major terms in the model that describe the angular distribution of microfacets, how many are visible from each angle, and how light reflects from each facet.

The first term D in the Torrance–Sparrow model describes the density of facets facing in any possible direction:

$$D = \frac{e^{-(\tan^2 \theta_h)/m^2}}{4m^2 \cos^4 \theta_h}, \quad (1.16)$$

where θ_h is the angle between the halfway vector \mathbf{h} and the surface normal \mathbf{n} . Notice that part of this term resembles a Gaussian, and this is not a coincidence: the Torrance–Sparrow model makes the assumption that the microfacet normals have a Gaussian distribution controlled

by a “roughness” parameter m . The $\cos^4\theta_h$ term occurring here is a change-of-basis term: it is included to properly normalize a probability distribution expressed in terms of the halfway vector.

The next term G in the Torrance–Sparrow model accounts for the fact that not all facets are visible from all directions, because they are hidden by the facets in front of them. It includes both “shadowing” and “masking” effects, representing occlusion from the point of view of the light and viewer, respectively:

$$G = \min \left\{ 1, \frac{2 \cos\theta_h \cos\theta_i}{\cos\theta_d}, \frac{2 \cos\theta_h \cos\theta_o}{\cos\theta_d} \right\}. \quad (1.17)$$

This formula is derived by considering a particular microgeometry: the microfacets are assumed to form V-shaped grooves in the surface, which are symmetric about the (macroscopic) surface normal.

Finally, the reflection from each facet is described by the Fresnel term F , which predicts that reflection increases toward grazing angles. This term arises from a solution to Maxwell’s equations on a surface:

$$F = \frac{1}{2} (F_{\perp} + F_{\parallel}) = \frac{1}{2} \left[\left(\frac{\sin(\theta_t - \theta_d)}{\sin(\theta_t + \theta_d)} \right)^2 + \left(\frac{\tan(\theta_d - \theta_t)}{\tan(\theta_d + \theta_t)} \right)^2 \right], \quad (1.18)$$

where $\theta_t = \sin^{-1}((\sin\theta_d)/\eta)$, η is the index of refraction of the surface, and the two terms represent the portion of reflected light polarized perpendicular and parallel to the plane of incidence. Note that this term involves the “difference angle” θ_d , as defined in Figure 1.7, which is the angle of incidence (and exitance) on a microfacet oriented to produce mirror reflection between the desired angles of incidence and reflection.

More recently, Ashikhmin et al. [3] generalized these types of microfacet BRDFs to allow expressing arbitrary half-angle distributions. They demonstrate how to modify these BRDFs to replace the analytic distribution in (1.16) with alternative analytic forms or tabulated (sampled) functions that can express arbitrary patterns.

More complex analytic BRDFs: In addition to models for primarily-specular surfaces, physically-based BRDFs have been derived

for rough diffuse surfaces (the Oren–Nayar model [146]), and for dusty surfaces (the Hapke/Lommel–Seeliger model, developed to model lunar reflectance [66]). They range in complexity from simple formulas that ignore many real-world effects to complex models that attempt to account for most actually observed surface phenomena (e.g., the He–Torrance–Sillion–Greenberg model [71]). While a detailed description of these models is beyond the scope of this study, they are frequently used in photo-realistic rendering systems. One drawback of these models, however, is that their additional complexity and many parameters can make it difficult or unstable to fit them to measured data.

Beyond Analytic BRDFs: Although we could continue to develop mathematical BRDF formulas of increasing sophistication that explains a greater variety of optical phenomena, over the past decade it has become increasingly practical to simply measure the BRDFs of real material samples [122]. In fact, this is one of the main theses of the avenue of research surveyed in this study: that measured data can capture a greater variety of real-world optical phenomena with greater accuracy than is possible with analytic models.

1.3 6D Datasets: SVBRDFs, BTFs, and Distant-Light Reflectance Fields

Of course, the BRDF is merely the beginning of our study of the appearance of materials. Real-world objects will exhibit more complex behaviors, such as a BRDF that changes from point to point on the surface. Adding two spatial dimensions to the four directional dimensions of the BRDF leads us to the six-dimensional Spatially-Varying BRDF:

$$SVBRDF(x, y, \theta_i, \varphi_i, \theta_o, \varphi_o). \quad [\text{sr}^{-1}] \quad (1.19)$$

Section 3 of this article describes the challenges of capturing, representing, editing, and analyzing these complex functions [94].

The study of the Spatially-Varying Bidirectional Reflectance-Distribution Function (SVBRDF) necessarily represents a shift from thinking of the appearance of “materials” to that of “objects,” and therefore requires considering the role of object geometry. In cases in

which the geometry is known (either because it is planar or because it has been scanned or modeled), the spatial dimensions of the SVBRDF are simply represented on a parameterization of that geometry. The SVBRDF is thus defined very close to the interface between a surface and the surrounding air, and it seeks to describe the scattering effects that occur at and immediately below this interface.

In many cases, however, it is impossible, difficult, or undesirable to model the scene geometry and to compensate for its effects during measurement. In this case, a 6D function may still be defined, with the spatial dimensions represented relative to some reference surface, or simply as image coordinates in a camera that was used for capture. In this case, the function is often called a *non-local reflectance field*, or simply *6D reflectance field*. It may be thought of as representing the *apparent* exitant light field $L(x, y, \theta_o, \phi_o)$ [105, 60] due to all possible directions of (non-local) incident light or, equivalently, the (4D) reflectance field $R(x, y, \theta_i, \phi_i)$ [30] for all possible viewing directions. When the reference geometry is planar, the term *Bidirectional Texture Function* (BTF) is frequently used [24].

BTFs and 6D reflectance fields, if sampled sufficiently densely, can represent non-local effects including those of foreshortening, occlusion, shadowing, refraction, subsurface and volumetric scattering, and inter-reflection. They are useful for objects that have significant mesostructure, or geometric structure that exists at or near the measurement scale. However, they give up the property that 4D “slices” at individual locations on the reference surface are proper physically-plausible BRDFs. Because of the non-local effects on apparent reflectance, they may fail to satisfy reciprocity or energy conservation and, even for surfaces with an underlying isotropic material, might not exhibit isotropy.

1.4 Subsurface Scattering

The SVBRDF and the BTF are not enough to characterize all materials. Many surfaces exhibit translucency: a phenomenon in which light enters the object, is reflected inside the material, and subsequently re-emerges from a different point on the surface. Such subsurface scattering can have a dramatic effect on appearance, as can be seen from

these computer graphics simulations that differ in only one respect: the left image simulates surface reflection only, while the right image includes subsurface scattering [150].

In order to cope with subsurface scattering, we will need to examine more complex appearance functions: those that can include the phenomenon of light leaving the surface at a different point than the one at which it entered.

1.4.1 The BSSRDF

The relevant function is known as the Bidirectional Scattering-Surface Reflectance-Distribution Function, or BSSRDF:

$$S(x_i, y_i, \theta_i, \varphi_i, x_o, y_o, \theta_o, \varphi_o) = \frac{dL(x_o, y_o, \theta_o, \varphi_o)}{d\Phi(x_i, y_i, \theta_i, \varphi_i)}. \quad [\text{m}^{-2} \cdot \text{sr}^{-1}] \quad (1.20)$$

This takes the SVBRDF and adds two more variables, representing the surface location at which the light leaves the surface: we are now up to a function of eight variables.

Unlike the BRDF, which is defined relative to input power averaged over a differential area, the BSSRDF is defined relative to input power at a single point. For this reason, the BSSRDF is expressed as a fraction of incident flux instead of incident irradiance, and its units are inverse squared meters times inverse steradians [142].

As we will see later in this article, the high dimensionality of this function leads to great difficulty in capturing and working with the BSSRDF directly, especially if a high sampling rate in each dimension is desired [57].

1.4.2 The Dipole Model

Because of the enormous size of the BSSRDF, approximations to it have become quite popular. One of the most powerful relies on the fact that, in many cases, the appearance is dominated by light that has reflected many times within the material. In this case, the details of each scattering event become unimportant, and the appearance is well approximated by thinking of light “diffusing” away from the location at which it enters the surface, much as heat might spread [83].

It turns out that the pattern of diffusion is well approximated by a dipole: a combination of a point light some distance *below* the point at which light entered the surface, and a *negative* light source some (slightly larger) distance above the surface. Combining the contributions of these two light sources with transmissive Fresnel terms for light entering and leaving the surface (which are simply one minus the reflective Fresnel equation given in (1.18)) yield a simple, yet powerful, model:

$$S = \frac{1}{\pi} F_t(\theta_i) R_d(\|\mathbf{x}_i - \mathbf{x}_o\|) F_t(\theta_o). \quad [\text{m}^{-2} \cdot \text{sr}^{-1}] \quad (1.21)$$

Because of the symmetry of diffusion, the model is effectively a function of only one variable: the distance between the points of incidence and exitance.

This dipole model, originally introduced in 2001, has become popular for simulating subsurface scattering in many materials, and we will see applications of it later in this article.

1.4.3 Homogeneous and Heterogeneous Scattering

Of course, the dipole approximation assumes a uniform material: the same amount of scattering everywhere on the surface. For more realistic surfaces, it might be necessary to add some of the complexity of the BSSRDF back in, by considering spatial variation. For example, in Figure 1.8 it is clearly visible that internal structure affects the scattering.

1.5 8D Reflectance Fields

While the BSSRDF is most often associated with subsurface scattering in translucent media, measured on the material/air interface, its definition is general enough to represent light transport from incident to exitant light rays in arbitrary configurations (e.g., the two rays do not even need to intersect). When thought of in this way, it is typical to refer to this function as an *8D reflectance field*, representing the full exitant 4D light field for each possible ray of a 4D incident light field. As with the 6D reflectance field, it is common for this function to be parameterized over some arbitrary reference surface enclosing the scene. Thus, an



Fig. 1.8 *Left*: a synthesized image with surface reflection only. *Right*: the same model with a simulation of subsurface scattering.

8D reflectance field, if sufficiently densely sampled, can describe scattering in participating media, interreflection among multiple objects, refractions, caustics, cast shadows, etc.

A major difference between 6D and 8D reflectance fields is that the latter can include effects due to local lighting. In contrast, the 6D reflectance field carries the implicit assumption of distant (directional or environment) lighting, and thus includes the net effect of equal illumination along all incident rays having the same direction. Thus, only the 8D reflectance field can capture, for example, the effect of an individual light ray on the surface, which is often of interest for translucent materials.

1.6 Generalizing Reflectance and Scattering

It turns out that even BSSRDFs and 8D reflectance fields do not cover all possible aspects of surface appearance. First, one could consider all of the functions discussed above as being dependent on the wavelength of light. Moreover, some surfaces are *fluorescent*: they emit light at different wavelengths than those present in the incident light. Other surfaces may have appearance that changes over time because of chemical changes, physical processes such as drying, or weathering.

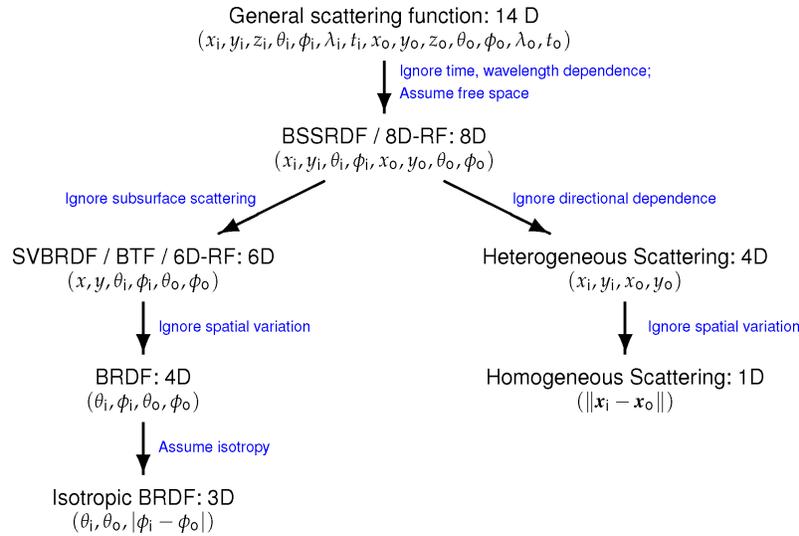


Fig. 1.9 Taxonomy of scattering and reflectance functions.

Still other surfaces might capture light and re-emit it later, leading to *phosphorescence* and other such phenomena.

Thus, a complete description of light scattering needs to add at least two wavelength and two time dimensions to the BSSRDF. Moreover, representing volumetric scattering adds two additional spatial dimensions, since this violates the assumption that radiance along light rays is constant.

So, we can think of all of the functions we have seen as specializations of a 14-dimensional scattering function, representing the distribution of light at arbitrary time, wavelength, position, and angle that is due to an incident light ray at some time, wavelength, position, and angle. While nobody has really tried to capture this full function, many efforts exist to capture one or more of its low-dimensional subsets (Figure 1.9). In fact, it can be argued that over the past decade, researchers have explored most of the subsets that “make sense,” up to the limits of acquisition devices.

2

Principles of Acquisition

Measuring reflectance requires the ability to strike a surface with visible light and sense the radiance it emits. Traditionally, this has been accomplished with a gonireflectometer. Consisting of a moving light source and a single moving photo-detector, this device is slow but provides highly repeatable Bidirectional Reflectance-Distribution Function (BRDF) measurements. Over the past two decades, significant effort has been devoted to streamlining the measurement process using cameras, light arrays, projectors, and generalized optics. Among other things, this has enabled the collection of large BRDF databases and the efficient measurement of higher dimensional scattering functions.

This chapter provides an overview of the acquisition process. We begin with the basic building blocks: lights sources and projectors to emit light rays; lenses and mirrors to bend them; and digital cameras to measure radiance that is reflected from a material. We then provide an overview of acquisition system designs for reflectance functions defined over domains of increasing dimension.

2.1 Traditional BRDF Measurement: The Gonioreflectometer

Consider the BRDF to be a function of five dimensions — four angular dimensions and one spectral dimension. Ignore, for now, its dependence on polarization and time, and do not allow fluorescence. One generally measures such a BRDF by illuminating a (locally) flat surface with a narrow beam of light with directed solid angle (ω_i) and placing a sensor to subtend an output directed solid angle (ω_o). The system is designed so that the input and output directions are known relative to the local coordinate system of the planar surface patch.

The classic BRDF measurement device is the four-axis gonioreflectometer. This device uses a combination of servo motors to position a source and photo-detector at various locations on a hemisphere above a planar material sample. The sensor is typically linked to a spectroradiometer or another optical assembly that permits recording of dense spectral measurements for each angular configuration of the source/sensor pair (e.g., [197]).

This measurement process is a lengthy one, and it can require days to measure a single material. The advantage of this approach, however, is that the system can be carefully calibrated, so that measurements become repeatable. It also provides the ability to capture dense spectral information, which is a tremendous advantage over most camera-based systems.

Acquisition time and equipment cost can be reduced if one is willing to restrict attention to isotropic BRDF. In this case, the angular domain is only three-dimensional, and one requires only three degrees of freedom (3DOF) in the acquisition system. For a gonioreflectometer built at Cornell University [108], this is accomplished by having two degrees of freedom (2DOF) in the orientation of the planar sample and one degree of freedom (1DOF) in the angular position of the source. Using this system, one can acquire 31 spectral samples per source/sensor configuration (10 nm increments over the visible spectrum); and capturing 1000 angular samples — a very sparse sampling of the 3D angular domain — takes approximately 10 hours.

For repeatability, any measurement of reflectance should be accompanied by a precise specification of the solid angles used to acquire

it, and precise guidelines for these specifications have existed for quite some time [142, 171]. Gonioreflectometers can be carefully constructed and calibrated so that these specifications are readily available, and for this reason, they are considered “gold standard” measurement devices for BRDF. But even with these devices, there are physical limitations that make accurate BRDF measurements difficult to obtain. Grazing angle effects, in which the incident and/or exitant directions are close to the measurement plane, are very hard to capture because of surface foreshortening; and the sensor’s noise level and limited dynamic range can corrupt measurements of strong specular reflections or very dim reflections.

2.1.1 The Importance of Scale

The BRDF is a derivative quantity, so any measurement device necessarily measures its average over finite intervals of space and solid angle. Indeed, “truly infinitesimal elements of solid angle do not contain measurable amounts of radiant flux” ([142], p. 7). We use the term *scale* in this chapter to refer to the size of the angular and spatial intervals used to measure a BRDF, and the importance of determining an appropriate scale cannot be overstated.

Consider the measurement geometry shown in Figure 2.1, where a portion of a planar sample is observed by a sensor through an optical

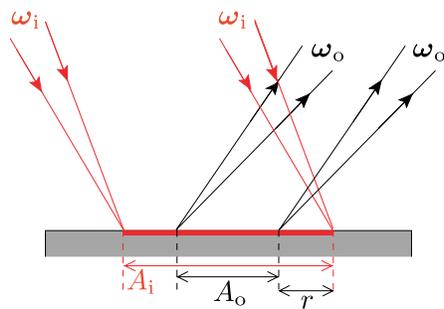


Fig. 2.1 BRDF measurement from a planar sample. Radiance emitted from surface patch A_o into solid angle ω_o is observed by the sensor (only extreme rays are drawn). This surface must be uniformly illuminated over larger area A_i , the scattering properties must be uniform across A_i , and subsurface scattering effects must be limited to a radius of less than r . (Adapted from [142], p. 25.)

system. The finite area of the sensor back-projects to a finite area on the surface A_o , and the sensor measures the radiant flux emitted from this area over a finite solid angle ω_o . The planar sample is illuminated by a bundle of light rays covering area $A_i \supset A_o$. Selecting a measurement *scale* requires making choices for the values of A_i , A_o , ω_i , and ω_o .

Precise BRDF measurements are obtained only when the finite solid angles ω_i and ω_o are small enough for the BRDF to be relatively constant over the directions within them. This means, for example, that measuring the BRDF of a perfect mirror requires a sensor with an infinitesimal aperture. Fortunately, for graphics and vision applications, sufficient precision can usually be obtained using “off the shelf” sensors that have small apertures and high sensitivity. Sufficient precision in ω_i is also quite achievable, but choosing appropriate values for the spatial intervals A_o and A_i is more difficult. In order to obtain a repeatable measurement, the surface scattering effects must be statistically uniform over A_i , where statistical uniformity means that any variations (e.g., due to interreflections, shadowing, and masking at a rough air/material interface) must occur over very small distances relative to the smaller area A_o . In addition, A_i must be large enough to guarantee that flux incident outside of A_i would not contribute significantly to the radiance reflected within A_o . This means that subsurface scattering cannot cause measurable fractions of incident flux to be transported over distances greater than r .

How do we verify that these spatial scale conditions are satisfied? One approach is to use reciprocity, which tells us that the BRDF must be unchanged when the directions of the source and sensor are exchanged. In theory, by physically exchanging the input and output directions in a measurement system (with the same angular and spatial intervals), one may verify this condition. Alternatively, one can gradually increase the areas A_o and A_i and determine when observable changes cease to occur in the estimated BRDF value. One can also measure the distance r by observing the reflection from a surface point intensely irradiated by a sharply focused beam, such as that from a laser (see Section 2.4). In this case, the radius of the band of diminishing radiance surrounding the intense central spot provides an

estimate of r . Finally, one may simply choose a spatial scale and then try to isolate the “direct” reflection component at that scale from the “non-direct” one. This separation can be approximately accomplished by modulating the incident light field, a procedure that is described in more detail in Section 2.5.2.

2.2 Image-based Acquisition

Vision and graphics applications often require representations of higher dimensional appearance functions (Spatially-varying BRDF (SVBRDF), Bidirectional Scattering-Surface Reflectance-Distribution Function (BSSRDF), etc.), and in order to measure these functions we require techniques that are more efficient than the gonireflectometer described in the previous section. To achieve this, we leverage arrays of light sources, digital projectors, and digital cameras to rapidly manipulate and sense visible light. A camera’s sensor array contains millions of photo-sensitive elements, and by using lenses and mirrors appropriately, these elements can be used to simultaneously collect a large number of reflectance samples. Similarly, the rays output by an array of sources or a digital projector can be directed onto a material sample from various directions with few moving parts.

What are we trading for the gain in efficiency obtained with these devices? Spectral resolution, for one. If a trichromatic sensor is used (e.g., Red, Green, Blue), one obtains only three spectral measurements for each angular configuration, and these measurements are weighted averages over large, overlapping intervals of the visible spectrum. This reduced spectral information is not sufficient for physically-accurate color reproduction, and it limits one’s ability to predict a material’s appearance under changes in source spectrum. Another limitation is the complexity of the required calibration procedure. In any camera-based system, one must estimate the geometric mapping from pixels in the camera (points in a plane) to rays in a three-dimensional coordinate system attached to the sample. To be precise, one must also estimate the exitant solid angle subtended by the camera’s aperture, and the spatial area A_o sensed by each pixel. Finally, one must estimate the radiometric mapping between the intensities recorded by a camera and scene

radiance values, including corrections for the vignetting in the camera's optical system and the camera's radiometric response function.

The complexity and fragility of the required calibration reduces the accuracy and repeatability of the measurements, but the significant efficiency gains often make this trade worthwhile.

When using a camera-based system for reflectometry, one must pay particular attention to scale. As depicted in Figure 2.1, it is essential that the spatial intervals A_o observed by each pixel are large enough to satisfy the requirement for statistical uniformity in the material. In general, this means that high-resolution images must be appropriately downsampled to obtain valid BRDF measurements.

2.2.1 Camera Calibration

As described above, a camera must be calibrated both geometrically and radiometrically to be used for reflectometry. The purpose of the former is to recover the mapping from the camera's pixels to rays in a three-dimensional coordinate system attached to a material sample. The purpose of the latter is to recover the mapping from each pixel's intensity value to a (relative) scene radiance value.

Geometric calibration involves the recovery of a camera's extrinsic parameters (position and orientation relative to a world coordinate system) and intrinsic parameters (focal length, radial distortion parameters, etc.). Free and reliable tools for geometric camera calibration are readily available [145, 12]. For a small number of cameras with similar fields of view, the most practical procedure uses multiple images of a planar checker-board pattern [210]. For camera arrays with a wider set of views, synchronous imaging of a small moving point light source, such as a light-emitting diode (LED), is a better alternative [27]. Note that none of these techniques provide information about the solid angle subtended by the camera's aperture, and in practice this is often left unspecified in camera-based reflectometry.

Radiometric camera calibration involves two stages. First, one must determine the radiometric response function of the camera. This is the non-linear mapping that often exists between the irradiance incident on the image plane and the recorded intensity. Standard methods for

doing this exist, and they usually involve imaging a calibration target or capturing multiple exposures of a static scene (see [61] for a recent review). A second step is required to recover the optical fall-off of the lens system. An imaging system with an ideal thin lens exhibits a relationship between scene radiance and image irradiance that falls-off as $\cos^4\alpha$, where α is the angle between the incoming chief ray and the optical axis. In a real camera, the optical fall-off includes vignetting effects and must be measured for each zoom and aperture setting. It can be measured, for example, by capturing an image of a cloudy sky through a diffuser. One then compensates for optical fall-off by dividing subsequent measurement images by this “vignetting image.”

A severe limitation that comes with the use of digital cameras for reflectometry is their small dynamic range. In order to measure high radiance values at specular peaks while maintaining sufficient signal-to-noise ratios in darker regions, one must acquire several images with different exposures and merge them into a single high-dynamic-range (HDR) image [31, 131]. This task is relatively straightforward once the radiometric camera response is known.

Another limitation of cameras is the finite depth of field induced by their optical designs. Only one fronto-parallel scene plane can be in perfect focus at any given time, and the depth of field (i.e., the range of depth values that are sufficiently sharp) is inversely proportional to the camera’s aperture. When an object is close to a camera and has significant volume, one often requires a small aperture setting to obtain a sufficient depth of field, and this, in turn, means that exposure time (and sensor noise) increases. The correct balance between depth of field and exposure time for a given acquisition system depends on the objects being considered, the amount of light available, and the noise levels of the cameras being used. This balance is determined largely through trial and error.

2.2.2 Light Source Calibration

Individual light sources are often used to create incident light fields for appearance acquisition, and these light fields are typically modeled in one of two ways. They are either modeled as parallel rays at a particular

orientation, as created by far-field source; or they are modeled as a set of rays that intersect at one point, as emitted by an ideal near-field point source. In either case, calibration consists of estimating the radiometric and geometric parameters of the light field model.

The parameters of a light model are determined from measurements collected by a geometrically and radiometrically calibrated camera. Geometrically, the position of a near-field source or the direction of a far-field source is recovered by placing shiny spheres in the scene, estimating their positions relative to the calibrated camera, and observing the points of mirror reflection (e.g., [100]). Radiometrically, one recovers the relative radiance along each ray (most sources do not emit uniformly), by imaging a planar diffuse target (made of a material such as Spectralon®) whose position is known relative to the light source and camera. Each pixel provides a measurement of the relative radiance emitted from a point on the planar target, and from this one can compute the relative radiance *incident* at that point by: (1) back-projecting the pixel ray to its corresponding 3D point on the target; (2) computing the surface normal and local light direction at that point; and (3) dividing by the scalar product of the normal and light vector to compensate for surface foreshortening. The result of this calculation is the relative radiance along one source ray.

In some cases, linear light sources and area light sources have proven effective, in which cases alternative calibration procedures are required (e.g., [48]).

2.2.3 Projector Calibration

Programmable light sources are those that allow the relative radiance along output rays to be independently controlled. Digital projectors are the most readily available programmable sources, and by replacing standard light sources with projectors in an acquisition system, one obtains greater flexibility in exchange for increased complexity and cost.

A projector is modeled exactly like a camera, so as described above, its calibration requires estimating the mapping between projector pixel intensity values and relative radiance values along corresponding rays

in 3D space. To achieve this calibration, one first calibrates a camera radiometrically and geometrically, and then uses this camera to make measurements of the projector's output [156]. To obtain the intrinsic and extrinsic parameters that describe the projector's geometry, a simple technique is to project a fixed checker-board pattern onto a planar surface that is placed at multiple orientations in the scene. By collecting images of these illuminated planes with a fixed camera, the intrinsic and extrinsic parameters of both the camera and projector can be recovered in the same manner as is done for a camera pair [12].

Radiometrically, a projector can be calibrated from as few as two images. First, a "vignetting image" for the projector is acquired by projecting an all-white pattern onto a planar diffuse target, capturing an image with a calibrated camera, and then estimating the relative radiance of each projector ray as described above. Then, one projects a grayscale "step" pattern (one square for each gray level) onto the same plane, records an image, converts this to radiance values along the projector rays using the known geometry, and then corrects for optical fall-off using the vignetting image. This yields the projector's non-linearity in tabular form.

An important disadvantage of using digital projectors is that they exhibit a variety of phenomena that are undesirable from the standpoint of appearance acquisition [209]. They often vibrate at a high frequency; their output radiance varies slightly over time; there is "light leakage" from one pixel to the next; the output radiance from an "all black" image is non-zero; and "pixelation" is augmented by dead space that exists between pixels. In addition, because of their optical design and relatively large aperture, projectors typically have a small depth of field.

Only some of these effects can be corrected. The effective "black level" of the projector, meaning the scene radiance it induces when projecting an all black image onto a particular scene, can be measured and then subtracted from subsequent measurements obtained with "patterned" projections onto the same scene. Also, even though pixelation and blurring cannot be completely eliminated, they can be reduced using signal processing techniques [209].

2.2.4 Colorimetric Calibration

As described above, camera-based acquisition systems typically provide only coarse spectral appearance information. If a single light source and trichromatic camera are used, then one obtains three spectral measurements of reflectance, each of which represents a weighted integral over the visible spectrum. This is ideally obtained using separate color filters and a monochrome camera, but for efficiency reasons, it is sometimes obtained by demosaicking a single image captured through a color filter array.

The International Color Consortium (ICC) maintains a standardized color management system that can be used to approximately translate measurements obtained with a particular camera under a particular light source to a standard colorimetric representation — usually CIE XYZ — that can, in turn, be appropriately transformed for any calibrated display. The translation from camera tristimulus values (Red, Green, Blue) to CIE XYZ is approximate because a camera’s spectral filters are not exact linear combinations of the CIE standard observer’s. (This means, for example, that the mapping from camera tristimulus values to XYZ is not one-to-one.) A way of obtaining a good approximate translation is to use established software tools that can create an “ICC profile” for any source/camera pair. A discussion of this procedure in the context of appearance acquisition is provided by Goesele et al. [56], who also describe how to combine ICC color management with HDR imaging.

This procedure creates an appearance representation that is valid for the illuminant spectrum used for acquisition. Appearance under changes in illuminant spectrum can only be approximated, and the usual method for doing so is through *generalized diagonal transforms* [40], meaning independent per-channel gains in a color space that is linearly related to XYZ. Based on perceptual experiments, a preferred color space for this form of “spectral re-lighting” is one known as the Bradford transform (see [38]), and this has been incorporated into the ICC color management system [169]. Alternatively, one can compute an optimized color space for relighting if one is given a set of illuminant spectra, spectral BRDFs, and per-channel sensor/display spectral distributions [40, 16].

When a projector is used as a light source, the measurement procedure is complicated by the existence of its three spectral filters. DLP projectors use a rotating color wheel to project colors, so in order to make an accurate appearance measurement under such a projector’s “white,” one should set the camera’s exposure time to be an even multiple of the filter wheel’s rotational period, or remove the filter wheel altogether.

Colorimetric calibration is especially important when multiple cameras and multiple light sources (or projectors) are used. In this case, any differences in the spectral characteristics of the devices can manifest themselves as noise in the aggregated reflectance measurements. To compensate for these differences, one must create separate ICC profiles for each camera/source pair and map measurements from each to a standard color space as described above.

2.2.5 Object Shape

If the materials being measured exist on arbitrarily-shaped surfaces, the shape of these surfaces must be known in the same coordinate system defined by the geometric calibration of the cameras and sources. From the perspective of measuring reflectance, three-dimensional surface reconstruction and alignment can thus be viewed as another calibration step. In this context, one of the requirements of the reconstruction system is that it recovers shape in a manner that is not biased by the material properties of the surface being measured. Indeed, having the calibration of the measurement system depend on the signal being measured is very inconvenient.

A common approach to recovering shape for general materials is to use structured lighting from lasers or projectors [7, 79], possibly in conjunction with stereo cameras [29, 208, 28]. For the purposes of appearance capture, one disadvantage of these approaches is that the recovered shape must be aligned with the images used for reflectometry, causing alignment errors that are manifested in noisy reflectance samples. Another disadvantage is that they do not directly estimate surface normals, which are ultimately required for measuring reflectance. Estimating surface normals from a range-scan requires differentiation of the discrete shape, and this is an additional source of noise.

By directly estimating surface normals from the same images being used for reflectometry, these two sources of noise can be reduced. Surface normals can be estimated using photometric stereo, but in its classic formulation [203], this assumes Lambertian reflectance and violates the requirement of being independent of reflectance. Developing more general photometric stereo techniques is an active area of research, and [115, 59, 2, 77, 137] are some recent examples. The most practical method to date is perhaps that of Ma et al. [112], which uses a sequence of three gradient-illumination patterns induced on a sphere around an object. Assuming either perfectly diffuse or perfectly specular reflectance, the image captured under each illumination pattern provides one component of the surface normal, and by using polarization to separate the diffuse and specular reflectance components (see Section 2.5.1), one can use this technique to recover high-quality surface normals at high frame rates.

Surface normal estimates that are truly independent from surface reflectance can be obtained by exploiting Helmholtz reciprocity. This approach relies on the acquisition of reciprocal pairs of images — images for which the positions of the light source and camera are swapped [211]. The disadvantage of this approach is that it requires correspondence between images captured from different viewpoints to compute surface normals. Practically, this means that the resolution of the recovered normal field may be lower than that of the input images. Photometric stereo, on the other hand, provides per-pixel normals whenever the underlying reflectance assumptions are satisfied.

Current best practice is to recover coarse geometry using structured lighting and then combine this with surface normal estimates obtained photometrically from the same images being used for reflectometry (e.g., [112, 195]). This approach leverages the fact that range scanners provide accurate low-frequency geometry while photometric methods provide good high-frequency detail. A convenient algorithm for combining these two sources of data has been developed by Nehab et al. [140]. It is also conceivable that accurate low-frequency shape could be obtained without active lighting using multi-view stereo techniques [165].

2.3 Acquiring BRDF, SVBRDF and 6D Reflectance Fields

Cameras, light sources, projectors, and optics can be combined in a variety of ways to measure surface reflectance, and this section provides a survey of some existing acquisition system designs. We consider appearance functions defined on domains having between three and six spatio-angular dimensions, from isotropic BRDFs to six-dimensional Bidirectional Texture Functions (BTFs) and spatially-varying BRDFs. The acquisition of eight-dimensional functions (BSSRDFs and 8D reflectance fields) is discussed in Section 2.4.

In designing an acquisition system, there are four competing factors that need to be considered: acquisition time; precision; cost; and material diversity, meaning the breadth of materials that are to be considered. It is possible to build an efficient system for measuring BRDF using spherical material samples, for example, but not every material can be “painted on” to a sphere.

Figure 2.2 includes six different systems that have been used to measure BRDF, BTF, and SVBRDF. One of the earliest camera-based

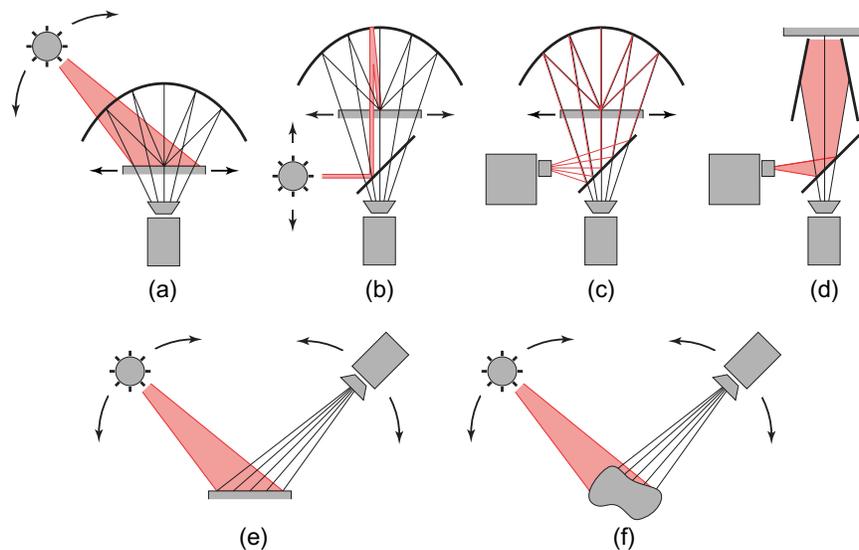


Fig. 2.2 Designs for acquiring BRDF, SVBRDF, and BTF (6D reflectance fields). In each case, a single image captured by a camera provides a dense sampling of a 2D slice of the spatio-angular domain. Table 2.1 provides corresponding references to the literature.

Table 2.1 Sample references corresponding to the designs in Figure 2.2. This is not a comprehensive list.

	BRDF	SVBRDF	BTF
(a)	Ward [191]	Extension to [191]	Extension to [191]
(b)	Dana and Wang [25]	Dana and Wang [25]	N/A
(c)	Ghosh et al. [54]	Extension to [54]	N/A
(d)	Inefficient	Han and Perlin [64]	Han and Perlin [64]
(e)	Inefficient	Dana et al. [24], Malzbender et al. [116], McAllister [126]	Dana et al. [24], Malzbender et al. [116]
(f)	Marschner et al. [119], Matusik et al. [122], Ngan et al. [141]	Stanford gantry [104] Debevec et al. [30]	Stanford gantry [104] Debevec et al. [30]

systems for measuring BRDFs was designed by Ward [191] and is shown in Figure 2.2(a). In this system, the radiance emitted by a planar sample is reflected from a half-silvered hemispherical mirror and captured by a camera through a fish-eye lens. In this way, a single image provides a dense sampling of the entire hemisphere of output directions, and by illuminating a homogeneous planar sample with light source subtending solid angle ω_i , this yields a densely sampled 2D slice of the BRDF: $f_r(\omega_i, \cdot)$. For isotropic materials, only one DOF in the light source is required, and this can be affected using a 1DOF rotation stage couple to the source arm. For anisotropic surfaces, the additional DOF is most easily achieved by coupling an additional 1DOF rotation stage to the material sample. Also, this system could be extended to handle SVBRDF by adding a 2D translation stage to the material sample as depicted in the figure. In this case, each image provides a 2D slice $f_r(\mathbf{x}, \omega_i, \cdot)$. Using this system, Ward claimed that an anisotropic BRDF could be densely sampled in 10 minutes. A very nice property of this system is that it allows the measurement of retro-reflection directions, meaning those for which the incident and reflected directions are equal. This is not possible with the gonireflectometer described in Section 2.1.

An alternative design, shown in Figure 2.2(b), uses a light source with a much smaller spatial extent. This corresponds to a smaller A_i in Figure 2.1, but with $A_i \supset A_o$ still large enough to satisfy the requirements for valid BRDF measurements. The light from the source is reflected from a parabolic mirror, which is also used to reflect the

exitant rays. This design is due to Dana and Wang [25] and it has the advantage of enabling the measurement of complete anisotropic SVBRDF without mechanical rotation stages. The incident direction is controlled by *translations* of the source and material sample, as depicted in the figure, and one expects measurements obtained by this system to be highly repeatable. Note that because the spatial extent of the light source is smaller than that of Figure 2.2(a), Dana’s system is appropriate for SVBRDF but not well suited to measuring general 6D reflectance fields, where light transport within and above the material sample can occur over greater distances.

One way to reduce calibration requirements and improve repeatability is to develop acquisition systems with few moving parts. This can be accomplished by replacing the single moving point sources in Figures 2.2(a–b) with a large number of fixed sources that provide a sampling of the incident hemisphere of directions. In addition to improving repeatability, this provides more flexibility in the set of incident light patterns used to illuminate the sample. Instead of sequentially activating each incident ray, one can simultaneously activate many to produce patterns corresponding to spherical harmonics. The advantage of this approach is that it reduces the dynamic-range requirements for the sensor and increases the signal-to-noise ratio. Schechner et al. [164] discuss these benefits and propose using a Hadamard basis to sample the space of illuminations from a discrete set of point sources.

Figure 2.2(c) achieves a very dense sampling of the incident hemisphere without motion using a digital projector. This design is due to Ghosh et al. [54], who actually use a pair of curved, rotationally-symmetric mirrors instead of the single mirror shown here. Instead of sequentially activating individual incident rays, they use an orthogonal basis defined over a radially symmetric zone of the input hemisphere. In addition to reducing dynamic-range requirements, this also reduces aliasing in angular samples. Figure 2.2(c) also depicts a 2DOF translation stage that is coupled to the material sample. This is not part of Ghosh et al.’s system but is a simple extension that would provide the ability to measure SVBRDF. Like Figure 2.2(b), the limited spatial extent of the incident illumination would make the system unsuited for measuring general 6D reflectance fields.

An alternative motionless design that *is* suitable for measuring both SVBRDF and BTF is the kaleidoscope-based system of Han and Perlin [64]. This system, which is depicted in Figure 2.2(d), is built around a polygonal mirror that narrows toward the material sample, and it allows multiple views of the material sample to be acquired by a single camera. For example, the triangular kaleidoscope in [64] allows for 22 distinct and simultaneous views. By adding a beam splitter and projector to the system, the sample can be also be illuminated from multiple directions simply by setting appropriate regions of the projected image to white. The main functional difference between Han and Perlin’s system (Figure 2.2(d)) and that of Ghosh et al. (Figure 2.2(c)) is that the latter provides lower angular resolution in exchange for high spatial resolution of the sample. Also, since the entire planar material sample is illuminated, the kaleidoscope system is more suitable for BTF measurement.

A fifth design based on planar material samples is shown in Figure 2.2(e). It is much like a gonioreflectometer, except that the single sensor is replaced by a camera. For this reason it is sometimes called a *spatial gonioreflectometer*. In comparison to Figure 2.2(a), for which each image measures a slice $f_r(\mathbf{x}_i, \boldsymbol{\omega}_i, \cdot)$, the spatial gonioreflectometer provides slices $f_r(\cdot, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$. That is, it yields a dense sampling of the spatial domain. To measure isotropic SVBRDF, only three degrees of freedom are required in the angular positions of the incident and reflected directions. A three-axis spatial gonioreflectometer was built by McAllister [126], using a 2DOF rotation stage coupled to the material sample and a 1DOF rotation stage coupled to a source arm. For anisotropic SVBRDF and general 6D reflectance fields, one requires 4DOF for the input/output directions that can be achieved by moving a single source/camera pair [24] or densely covering the visible hemisphere with multiple cameras and sources (e.g., [116]). Note that while it can be used for measuring homogeneous materials (BRDF), the spatial gonioreflectometer is no more efficient than the traditional gonioreflectometer when the camera and source are distant from the sample (so that the input/output rays are parallel). To gain efficiency, a near-field camera and/or source must be used [85].

An alternative approach to using curved mirrors is to use a curved material sample instead of a planar one (Figure 2.2(f)). This permits the efficient measurement of BRDF as well as the measurement of SVBRDF and 6D reflectance fields that are defined over non-planar shapes. Efficient BRDF measurement is achieved when the curved surface has uniform material properties and is observed at an appropriate scale [119, 111]. Since the surface normal varies from point to point on the surface, a single image under directional lighting provides a very dense (near continuous) sampling of a two-dimensional slice of the BRDF domain. For isotropic materials, a sphere can be used with a fixed camera and a 1DOF rotation stage coupled to the source arm [122]. For anisotropic surfaces, one can use a cylinder and increase efficiency using multiple material strips that are cut at different orientations relative to the tangent direction [141]. The cylinder provides one degree of freedom in its surface normal, and two more degrees of freedom are obtained by rotating the cylinder and the source. The fourth and final DOF comes from the number of “strips,” which is typically coarsely sampled.

Finally, instead of using spheres or cylinders as curved material samples, one can use arbitrary shapes as long as the geometry is known in the coordinate systems of the camera and light source. This is important for live materials such as human skin [117, 120] and the skins of fruits that cannot be “painted on” to a plane, sphere, or cylinder.

Unlike measurement systems based on planar materials (Figures 2.2(a–e)), each dense 2D slice obtained from an image of a curved sample does not correspond to fixed input or output direction. Instead, the sampling pattern of a curved sample is best understood using an alternative BRDF parameterization, such as Marschner’s [117] or the halfway/difference parameterization of Rusinkiewicz [159] (discussed on page 86), that aligns the per-image 2D slices with the coordinate axes. For orthographic camera projection and distant lighting — or more generally, when scene relief is relatively small — a single image of a curved surface provides BRDF samples lying in a plane of constant difference angle (θ_d), since this angle is independent of the surface normal. In the special case of isotropic materials, this means that while each

(orthographic) image provides only one sample of the θ_d -dimension, it represents a nearly continuous sampling of θ_h and ϕ_d . As a result, a *set* of images provides dense sampling of (θ_h, ϕ_d) , but only as many samples of θ_d as there are images. Conveniently, this irregular sampling obtained from image-based data corresponds well with the behavior of general BRDFs, which vary slowly in the sparsely sampled θ_d -dimension, at least away from grazing angles [136]. At the same time, by imaging curved surfaces, one efficiently obtains high sampling rates of the half-angle θ_h , which are necessary to recover high-frequency variation (e.g., due to specular highlights) that is generally observed in this dimension.

The design in Figure 2.2(f) can also be used to measure SVBRDF and 6D reflectance fields. In order to capture the SVBRDF or 6D reflectance field of a curved surface, the acquisition system must be able to position sources and cameras over the entire sphere in order to sample the hemisphere of directions above every surface tangent plane. One device for accomplishing this task is a spherical gantry, the first example of which was built at Stanford University [104]. To get a sense of how many images are required to densely sample the reflectance field of a regular surface, we can perform a simple counting exercise. When the shape of the object is known and the source and view directions are given, each pixel in an image provides one sample of the reflectance field (or the BRDF at a particular surface point). Sampling the 4D angular domain at every surface point therefore requires capturing images from the full double sphere of view and illumination directions. Obtaining 5° or 1° angular sampling rates for these spheres requires millions or hundreds of millions of images, respectively.

To reduce measurement time is to capture the images more quickly using a system with an array of cameras that acquire in parallel and either a rapidly-moving source (e.g., [30]) or an array of sources that are activated sequentially (e.g., [116, 51, 196, 195]). This streamlines the acquisition process at the cost of an increase in calibration complexity (and fragility). For SVBRDF, another strategy for reducing the measurement burden is to decreasing the number of required images by exploiting symmetries (e.g., isotropy) and spatial regularities. This is discussed in more detail in Section 3.2.2 (page 136).

One very important consideration for reflectometry using curved surfaces is the presence of interreflections on non-convex shapes. When there is mutual illumination between distinct surface points, reflectance values can no longer be directly inferred from intensities observed on the image plane. For surface that are not mirror-like, mutual illumination can be removed from the images by modulating the input illumination with a moving binary pattern (see Section 2.5.2).

2.3.1 Passive Acquisition

All of the techniques described above rely on active illumination, meaning that illumination (and viewpoint) are carefully controlled during acquisition. An attractive alternative is to use passive techniques that recover reflectance information from one or more “natural” images via inverse rendering. Most inverse rendering problems are ill-posed, and in order to solve them, one must typically make strong assumptions about the class of materials that are present in the scene. (An early review of inverse rendering problems and techniques is provided by Patow and Pveyo [147].)

One approach is to use a parametric BRDF model, since these impose strong constraints on reflectance. One notable example is the work of Boivin and Gagalowicz [10], who employ the anisotropic Ward model, and show that when the scene geometry, camera position, and lighting are all known, one can estimate BRDF parameters from as little as one image. In their system, surfaces are manually grouped into regions of homogeneous reflectance, and then BRDF parameters are iteratively updated by comparing rendered results with the input image. In other work, parametric BRDF models have been shown to enable the simultaneous recovery of reflectance and shape (e.g., [50]), or reflectance and illumination (e.g., [67]).

Ramamoorthi and Hanrahan [155] derive an elegant framework for inverse rendering *without* parametric BRDF models by interpreting the rendering equation as a convolution. This yields an important theoretical tool that, among other things, enables the recovery of reflectance through de-convolution when the scene lighting and surface geometry are known, and when the complete 4D light field is observed. If, instead

of the entire light field, only a single image is given as input, one can apply the same approach, but in this case one is limited to so-called “radially-symmetric” BRDF, meaning those that are one-dimensional and symmetric about the reflection vector [155].

Developing inverse rendering techniques that are able to handle more general reflectance models, meaning those with are neither parametric nor radially-symmetric, remains an active area of research.

2.4 Acquiring BSSRDF

A general BSSRDF is a function of nine dimensions if we include the spectral dimension. Even if we ignore the dependency on wavelength, densely sampling an eight-dimensional space is an extremely burdensome process. To make it tractable, we reduce the dimensionality of the BSSRDF by expressing it in a parametric form. A number of parametric models have been explored. Each is targeted to a particular class of materials and each requires an acquisition system with particular capabilities.

The BSSRDF can be expressed as a linear combination of a single scattering term and a multiple scattering term. The former describes light that refracts below the surface and is scattered only once before refracting out [65], while the latter describes light that undergoes multiple scattering events below the surface interface. When the material is homogeneous, has a smooth surface, and is highly scattering, both terms can be described with high accuracy by a parametric model with relatively few degrees of freedom. When these conditions are satisfied, scattering within the material can be described by a diffusion process, which, in turn, can be approximated by a simple dipole model [83]. Accordingly, the BSSRDF can be written:

$$S = S^{(1)}(\mathbf{x}_o, \boldsymbol{\omega}_o, \mathbf{x}_i, \boldsymbol{\omega}_i) + \frac{1}{\pi} F(\boldsymbol{\omega}_i) R_d(\|\mathbf{x}_i - \mathbf{x}_o\|) F(\boldsymbol{\omega}_o), \quad (2.1)$$

where the first and second terms represent single and multiple scattering, respectively. The multiple scattering term is parameterized by only three spectrally-varying parameters. The first is the index of refraction (usually denoted by $\eta(\lambda)$). This is usually assumed to be constant over the visual spectrum, and it describes the Fresnel

transmittance functions, $F(\omega_i)$ and $F(\omega_o)$, that govern reflection and refraction at the (smooth) material/air interface. The analytic forms of these functions can be found, for example in [148]. The other two parameters are the *reduced scattering coefficient* $\sigma'_s(\lambda)$ and *absorption coefficient* $\sigma_a(\lambda)$, which together define an analytic expression for $R_d(r)$ in (2.1). (An alternative, more intuitive, parameterization uses the *total diffuse reflectance* and *translucency* or diffuse mean free path [82, 176].)

A variety of systems have been proposed for measuring the parameters of this model. All of them focus on the multiple scattering term and are based on the assumption that single scattering effects ($S^{(1)}$ in (2.1)) are negligible far from the point of incidence. This is true whenever the scattering albedo is close to one [82], meaning that the material modulates radiation predominantly through scattering as opposed to absorption (see [13]). Jensen et al. [83] measure the absorption coefficient and reduced scattering coefficient using a narrow focused beam of illumination and a camera (Figure 2.3). These measurements are made independent of the material’s index of refraction, and they provide three wide-band spectral measurements of each parameter. Weyrich et al. [195, 196] obtain similar measurements for human skin using a linear array of optical fibers that are couple to a camera (see Section 5.4.2.1, page 166). In both systems, the index of refraction is not measured and an assumed value (typically 1.3–1.5) is used instead. Also, neither of these systems measure the phase function

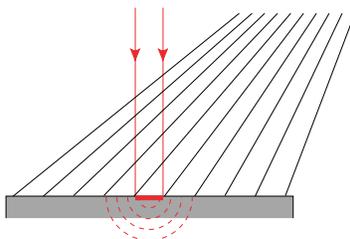


Fig. 2.3 Measuring a BSSRDF. A narrowly focused beam illuminates a homogeneous translucent material, and the radiance emitted from each surface point is recorded by a camera. When the air/material interface is smooth and the index of refraction is known, each pixel provides a measurement of the “subsurface function” $R(\mathbf{x}_i, \mathbf{x}_o)$.

of the material, which is necessary to compute the single scattering term of (2.1). This can be added to the model by using a parametric phase function (e.g., the Henyey–Greenstein phase function [74]) with a manually tuned parameter.

The model of (2.1) assumes the material/air interface to be smooth. This can be relaxed to account for microstructure (surface roughness) by adding a spatially varying BRDF and generalizing the Fresnel transmittance functions. If single scattering is ignored, one can write the BSSRDF as [34]:

$$S = \frac{1}{\pi} \rho(\mathbf{x}_i, \boldsymbol{\omega}_i) R_d(\|\mathbf{x}_i - \mathbf{x}_o\|) \rho(\mathbf{x}_o, \boldsymbol{\omega}_o), \quad (2.2)$$

where

$$\rho(\mathbf{x}, \boldsymbol{\omega}) = 1 - \int f_r(\mathbf{x}, \boldsymbol{\omega}, \boldsymbol{\omega}_i) \cos \theta_i d\boldsymbol{\omega}_i$$

is one minus the “directional-hemispherical reflectance factor” ([142], p. 12), or the fraction of incident radiant flux that is transmitted by the rough material/air interface. Note that the domain of integration in the expression above is the entire incident hemisphere.

So far, we have considered homogeneous materials with subsurface scattering properties that do not change from point to point. Most interesting materials are not homogeneous, however, and in practice it is necessary to adapt the models of (2.1) and (2.2) (and the acquisition system) to handle inhomogeneity. Most approaches do so by decomposing reflectance into local (BRDF) and non-local (BSSRDF) components (at a particular scale!), and then further factoring the non-local component. Accordingly, the emitted radiance is written as:

$$\begin{aligned} L(\mathbf{x}_o, \boldsymbol{\omega}_o) &= \int f_r(\mathbf{x}_o, \boldsymbol{\omega}_o, \boldsymbol{\omega}_i) L(\mathbf{x}_o, \boldsymbol{\omega}_i) \cos \theta_i d\boldsymbol{\omega}_i \\ &+ \iint f_i(\mathbf{x}_i, \boldsymbol{\omega}_i) R(\mathbf{x}_i, \mathbf{x}_o) f_o(\mathbf{x}_o, \boldsymbol{\omega}_o) L(\mathbf{x}_i, \boldsymbol{\omega}_i) \cos \theta_i d\boldsymbol{\omega}_i dA, \end{aligned} \quad (2.3)$$

where the spatially-varying BRDF f_r , the spatially-varying transmittance functions f_i and f_o , and the subsurface term $R(\mathbf{x}_i, \mathbf{x}_o)$ are either tabulated functions or low-parameter analytic models.

Acquiring the parameters of such a model, one requires an acquisition system that combines a spatial gonioreflectometer or spherical gantry (Figure 2.2(a,d,e)) with a subsurface measurement system. If acquisition time is an important consideration, the number of required measurements can be reduced by assuming a factored and/or parametric form for $R(\mathbf{x}_i, \mathbf{x}_o)$. Below are four examples:

- Jensen et al. [83] suggest using a dipole-diffusion model with reduced scattering coefficients and absorption coefficients that vary as a function of incident point on the material interface: $\sigma'_s(\mathbf{x}_i, \lambda)$ and $\sigma_a(\mathbf{x}_i, \lambda)$. This is a coarse approximation, of course, since it does not model volume effects, but the parameters can be measured very efficiently. Tariq et al. [176] use a single projector-camera pair to acquire spatially-dense estimates of $\sigma'_s(\mathbf{x}_i, \lambda)$ and $\sigma_a(\mathbf{x}_i, \lambda)$ of a human face in under a minute.
- Weyrich et al. [195] represent the non-local reflectance of faces using a dipole-diffusion model with fixed material parameters $\sigma'_s(\lambda)$ and $\sigma_a(\lambda)$ that is modulated by a spatially-varying scalar function: $R_d(\|\mathbf{x}_i - \mathbf{x}_o\|)M(\mathbf{x}_o, \lambda)$. They show that for human faces, the parameters of this subsurface model can be estimated directly from images acquired by a spherical gantry — the same images that are used to estimate the local (spatially-varying BRDF) component.
- Fuchs et al. [43] also use a scalar modulating function, but replace the two-parameter dipole-diffuse model with a linear combination of exponential functions:

$$R(\mathbf{x}_i, \mathbf{x}_o) = M(\mathbf{x}_o) \sum_k c_k(\mathbf{x}_i) e^{d_k(\mathbf{x}_i) \|\mathbf{x}_o - \mathbf{x}_i\|}. \quad (2.4)$$

The parameters of this model are measured at three wavelengths by illuminating a surface at many points sequentially using three optically-coupled lasers.

- Peers et al. [150] efficiently measure a non-parametric subsurface function $R(\mathbf{x}_i, \mathbf{x}_o)$ using a projector that illuminates multiple surface points in parallel. When these surface points are sufficiently far apart, a single image — with the Fresnel

transmittance effects removed — provides dense measurements of $R(\mathbf{x}_i, \cdot)$ for many \mathbf{x}_i .

- Unlike the three approaches above, which use parametric (Fresnel) transmittance functions and spatially-varying forms for $R(\mathbf{x}_i, \mathbf{x}_o)$, Tong et al. [179] handle inhomogeneity using measured, non-parametric, and spatially-varying representations of the transmittance functions f_i and f_o . This model is fit to measurements obtained by a gantry that combines a spatial gonireflectometer with lasers positioned at multiple input angles. It is a very flexible representation, capable of describing some general reflectance fields in addition to BSSRDF.

2.5 Separating Reflection Components

For non-conducting materials (dielectrics or non-metals), a typical image can be considered a linear combination of component images, with each component resulting from a distinct reflection mechanism. If we choose a particular measurement scale, then the radiance emitted from such a surface can be written as a sum of two terms,

$$L(\mathbf{x}_o, \boldsymbol{\omega}_o) = \int f_r(\mathbf{x}_o, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) L(\mathbf{x}_o, \boldsymbol{\omega}_i) \cos \theta_i d\boldsymbol{\omega}_i + \iint S(\mathbf{x}_i, \boldsymbol{\omega}_i, \mathbf{x}_o, \boldsymbol{\omega}_o) L(\mathbf{x}_i, \boldsymbol{\omega}_i) \cos \theta_i d\boldsymbol{\omega}_i dA, \quad (2.5)$$

corresponding to local reflection (BRDF) and non-local reflection (BSSRDF) at that particular scale. Likewise, the local BRDF component can be further factored into reflection that occurs exactly at the air/material interface and that which involves local subsurface effects.

These components differ in their dependence on wavelength, polarization, and the angular configuration of sources and sensors. For many objects, they also differ in how they change spatially across the surface. By separating images according to these reflectance components, their different behaviors can be exploited to reduce the acquisition time and the size of the resulting representation.

2.5.1 BRDF: Interface and Body Reflection

First, consider the case in which the measurement scale is chosen to be large relative to the subsurface scattering effects of a material (i.e., $A_o \gg r$ in Figure 2.1). In this case, there are no non-local effects, and the appearance of the material can be completely described by the first (BRDF) term in (2.5). We are interested in separating the appearance of this material into two components, corresponding to reflection at the air/material interface and reflections below the surface. These two components are termed the *interface* and *body* components [167] (*specular* and *diffuse* are also commonly used), and they can be separated using two physical properties: color and polarization.

The specular component of reflection corresponds to the portion of the incident radiant flux that is reflected at the surface boundary. For a perfectly smooth surface, this component is described by a “specular spike” occurring in the mirror direction, and the reflected flux can be computed using the Fresnel equations. Rough surfaces are more common, and in these cases, the specular component takes the form of a “specular lobe,” whose shape is determined by the surface microstructure (recall the Cook–Torrance model from Section 1). The diffuse reflection component is induced by flux that penetrates the surface and multiply refracts internally before refracting back out into the air. Diffuse-reflected flux is randomly polarized, and it is distributed more or less uniformly over the output hemisphere.

The diffuse and specular components differ in their dependence on wavelength, polarization, and the angular configuration of the input and output directions. For many objects, they also differ in how they change spatially across the surface, and this is shown in the example of Figure 2.4. In general, the SVBRDF $f_r(\mathbf{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ is a function of spatial (\mathbf{x}) and angular ($\boldsymbol{\omega}_i, \boldsymbol{\omega}_o$) dimensions. The diffuse component is well-represented by a Lambertian model for a wide range of input and output angles, but for surfaces with significant texture, it can change quite quickly from point to point. Said another way, it varies slowly in the angular dimensions, but varies quickly in the spatial dimensions. The specular component is often just the opposite. It tends to exhibit rapid changes in the angular dimensions, but often changes slowly over

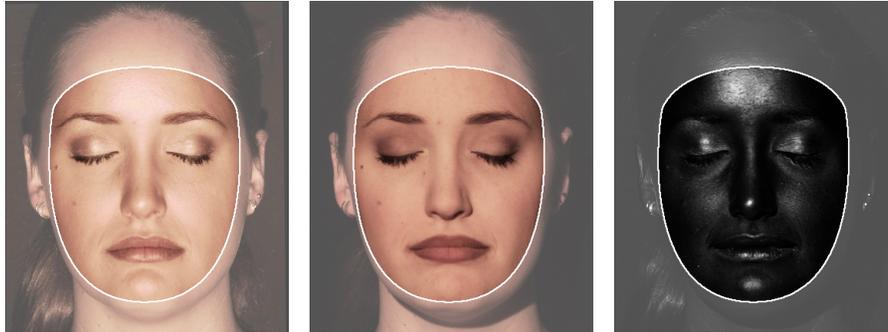


Fig. 2.4 An image that has been separated into diffuse (middle) and specular (right) components using polarization. The diffuse component is approximately Lambertian but varies rapidly between surface points. The specular component exhibits more complex variations in the angular domain, but varies smoothly (if at all) across the surface.

much of an object's surface. By separating the diffuse and specular contributions to a set of images of an object, these behaviors can be exploited to reduce the number of required input images [162, 212]. This will be discussed in more detail in the next chapter; here we focus on techniques for separating the two components in an image.

2.5.1.1 Polarization

The diffuse and specular components of an image are typically separated using color, polarization, or both. The polarization of light refers the orientations of the (orthogonal) planes of vibration for the electric and magnetic fields along the direction of propagation of an electromagnetic wave. A single ray can have linear, circular, or elliptical polarization, with each referring to the path traced out by the electric field vector as the wave propagates. We speak of collections of light rays as being, *polarized*, *unpolarized*, or *partially polarized*. Polarized light is coherent in terms of its polarization state, whereas unpolarized (or randomly polarized) light is completely incoherent. Partially polarized light refers to a mixture of randomly polarized and coherently polarized rays.

Before now, we have largely ignored polarization, and have focused instead on the angular and spatial distributions of incident and reflected

flux. Polarization reflectance effects can be quite complicated, and a completely general description might require as many as 16 BRDFs — a four-by-four matrix in terms of Stokes parameters — for a single material ([142], p. 32). Fortunately, polarization effects for non-conducting surfaces can be described and exploited quite simply.

The simplest way to separate specular and diffuse reflections is by “cross-polarization,” in which linear polarizing filters are placed on the source and camera (e.g., [128]). Most light sources emit unpolarized light, and by passing this through a linear polarizing filter, one obtains linearly polarized light at half the original radiance. When linearly polarized flux is incident on a non-conducting surface, the reflected flux will be partially linearly polarized. It consists of the diffuse component, which is unpolarized, and the specular component, which maintains a linear polarization. By placing another linear polarizing filter in front of the camera and orienting it with the polarization direction of the specularly-reflected flux, one obtains an image consisting of the complete specular component and half the diffuse component: $I_1 = I_d/2 + I_s$. Then, by rotating the camera’s filter by 90° , the specular component is blocked, and this second image contains the (attenuated) diffuse component only: $I_2 = I_d/2$. From these two images, the diffuse and specular components are simply given by $I_d = 2I_2$ and $I_s = I_1 - I_2$. Figure 2.4 shows an image that was decomposed in this way.

One potential complication in this process is that the polarization plane of the specularly-reflected flux depends on the material’s index of refraction (which may vary with wavelength) and the orientation of the incident polarization relative to the incident plane. This relationship is described by the Fresnel equations for a non-conductor, and by assuming an index of refraction (1.5, the refractive index of ordinary window glass, is the usual choice), one can design polarized lighting arrays with optimal polarization fields [112]. For a single source, however, the simplest strategy is to polarize the incident flux in the direction orthogonal to the plane spanned by the light direction and the camera’s optical axis. This way, the polarization of the specular-reflected flux will remain in that direction, regardless of the material’s index of refraction.

Another effect to be aware of is that diffuse reflectance can become partially polarized when the angle between the surface normal and view direction is large [200]. This means that polarization-based separation will be less accurate near the occluding contour of an observed object.

Finally, it is worth re-emphasizing that this strategy cannot be used for metallic materials or semiconductors, since reflectance in these cases is due purely to surface effects. The imaginary component of the index of refraction is non-zero in these cases, and the Fresnel equations that govern interface effects are more complicated (e.g., [148]). For example, linearly polarized incident flux leads to elliptically polarized reflected flux in general. The differences between the way in which different materials (especially conductors vs. dielectrics), and the way in which these vary geometrically, can be exploited in other ways than described here. This is especially useful in uncontrolled environments, where the polarization state of the illuminant cannot be controlled [200, 201].

2.5.1.2 Wavelength/Color

If we consider the wavelength dependence of the BRDF, it becomes a function of five dimensions, $f_r(\lambda, \omega_i, \omega_o)$. In many cases, this function can be factored into separated functions of wavelength and geometry:

$$f_r(\lambda, \omega_i, \omega_o) \approx g(\lambda)f(\omega_i, \omega_o). \quad (2.6)$$

This is typically a good approximation, but not always. For example, the underside of a compact disc exhibits joint angular and spectral variations caused by diffraction effects, and these cannot be described by this factored form.

As mentioned above, the BRDF of many materials, including all dielectrics, can be decomposed into two additive components: a specular (interface) component and a diffuse (body) component. Utilizing the spectral/angular factorization described above, the BRDF of such materials can be written as:

$$f_r(\lambda, \omega_i, \omega_o) = g_d(\lambda)f_d(\omega_i, \omega_o) + g_s(\lambda)\tilde{f}_s(\omega_i, \omega_o).$$

Additionally, since the index of refraction of many surfaces is constant over the visible spectrum [97] the function $g_s(\lambda)$ can be considered

constant, and this leads to the expression:

$$f_r(\lambda, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = g_d(\lambda) f_d(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) + f_s(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o), \quad (2.7)$$

where $f_s(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = g_s \tilde{f}_s(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$. In this expression, the function $g_d(\lambda)$ is often referred to as the *spectral reflectance* of the material.

Equation 2.7 is at the core of Shafer's *dichromatic model* of an image [167]. When such a surface is lit by a single illuminant and imaged by a trichromatic sensor, the observed colors can be written as linear combinations of a source color and diffuse color [167]. We can write the three channels of the observed image $\mathbf{I}_{RGB} = \{I_k\}_{k=R,G,B}$ as:

$$I_k = \sigma_d d_k + \sigma_s s_k, \quad (2.8)$$

where σ_d and σ_s are “geometric scale factors” that depend on the light positions and specular BRDF, and

$$d_k = \int e(\lambda) g_d(\lambda) c_k(\lambda) d\lambda \quad (2.9)$$

$$s_k = \int e(\lambda) c_k(\lambda) d\lambda. \quad (2.10)$$

Here, $e(\lambda)$ is the spectral power distribution of the incident illumination, $g_d(\lambda)$ is the spectral reflectance of the surface, and $c_k(\lambda)$ is the spectral sensitivity of a linear sensor. We define $\mathbf{d} = \{d_k\}_{k=R,G,B}$ and $\mathbf{s} = \{s_k\}_{k=R,G,B}$ to be the *diffuse color* and *source color* mentioned above. Empirically, this model has shown to be suitable for a diverse set of materials, including certain types of plant leaves, cloth, wood, and the skin of fruits [97, 178, 72].

The dichromatic model provides a means of using color for separating diffuse and specular reflection components. Unlike polarization, however, one cannot achieve this separation unambiguously using color. Even when the source color \mathbf{s} is known, there is a one-parameter family of diffuse colors that satisfy (2.8), and thus a one-parameter family of possible separations. To choose a good separation from the set of possibilities, the traditional approach is to exploit spatial coherence in the distribution of diffuse colors on a surface (e.g., [87, 114, 138, 167, 175]) and/or the distinct angular behavior of diffuse and specular effects [161].

2.5.2 Local and Non-local Reflection

In the previous section, we assumed that the measurement scale is chosen to be large relative to the subsurface scattering effects of a material. When this assumption is relaxed, one can consider three distinct reflection mechanisms instead of two. In addition to the specular (interface) component and local body component, there is a *non-local* body component that results from subsurface scattering over a radius that is significant relative to the measurement scale. Looking back at Figure 2.1, the difference between the local and non-local body components is that the former includes effects for which the radius r is much smaller than A_o , while the latter includes effects for which r is comparable to, or greater than, the measurement scale A_o .

Both the local and non-local body components exhibit the color and polarization properties described above, and both can be separated from the specular component using the same techniques. These physical properties cannot be used to separate the two body components from each other, however; for this we require another approach.

Nayar et al. [137] have recently shown that the local and non-local reflection components can be separated by modulating the illumination by high-frequency patterns. When a surface is illuminated by a high-frequency projected pattern, the exitant radiance due to local reflections exhibits rapid spatial variation, but since non-local reflections generally act as a low-pass filter, the exitant radiance due to this component is very smooth.

Consider the measurement geometry in Figure 2.5, where a rough inhomogeneous material is illuminated by directional lighting that is modulated by a binary pattern. As shown in Figure 2.1, the sensor records the emitted radiance over finite spatial and angular intervals. Suppose two measurements are taken, one with the binary pattern as shown, and one with the pattern inverted, and for reasons that will become clear in a moment, denote these images by I_{\max} and I_{\min} , respectively. If the spatial frequency of the binary pattern is high relative to the translucency of the material (i.e., its period is much smaller than the diffuse mean free path), these two measurements can

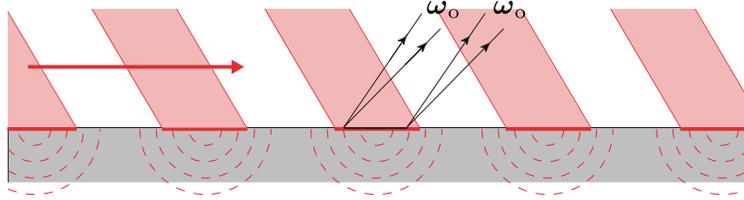


Fig. 2.5 Exitant radiance due to local and non-local reflections can be separated using high-frequency structured lighting. A checker-board illumination pattern is translated across the surface and the maximum and minimum intensities are recorded at each pixel of a camera. The local and non-local contributions (for this measurement scale) are trivially recovered from these measurements.

be written as:

$$I_{\max} = I_d + I_g/2$$

$$I_{\min} = I_g/2,$$

where I_d and I_g are the direct and global (or non-direct) components of the measurement that would be created if the binary pattern contained no zeros. In other words, given these two measurements, we can compute $I_d = I_{\max} - I_{\min}$ and $I_g = 2I_{\min}$, which provide good approximations to the two terms on the right of (2.3).

By illuminating a surface with a checker-board pattern and observing it with a camera, this same separation can be accomplished in parallel at every surface point. In practice, one shifts the checker-board pattern across the surface in order to obtain more accurate estimates of the maximum and minimum intensities at each pixel, and one then compensates for the fact that the radiance in dark regions of the binary pattern is not exactly zero [137]. Since other global effects, such as volumetric scattering and interreflections between surface points, also behave as low-pass filter of illumination, this same procedure can be used to isolate the direct component from these effects as well. (This can be leveraged for a variety of applications, including the reconstruction of translucent surfaces [15], and rapid estimation of scattering parameters for human faces [176].)

The definitions of local and non-local are scale-dependent, and will therefore vary with A_o and A_i in Figure 2.5. An illustration of this scale

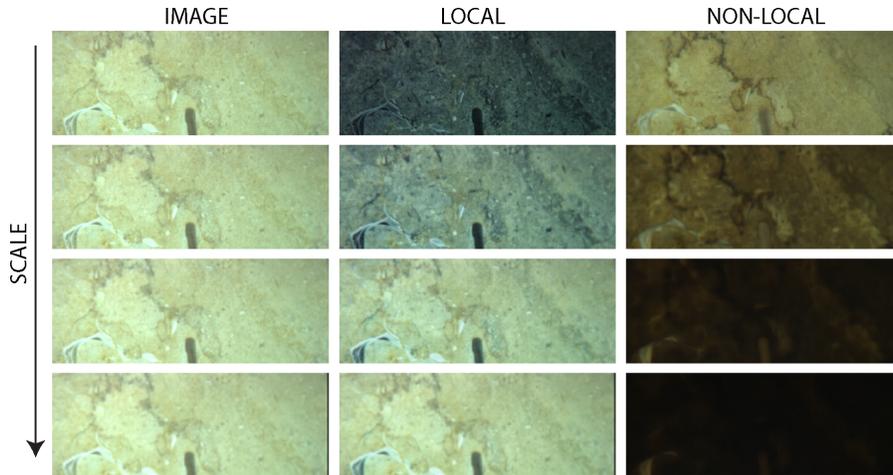


Fig. 2.6 The designation of reflection being local or non-local depends on measurement scale. Shown are the separated components using the method of Nayar et al. [137] for various resolutions of the image plane and incident checker-board pattern. The dominant mechanism of reflection goes from non-local to local as scale increases. (Figure courtesy of the Columbia University Computer Vision Lab (CAVE), <http://www.cs.columbia.edu/CAVE/>)

dependence is shown in Figure 2.6, where the separation is performed with A_i and A_o successively increased by a factor of two. At a small scale, the majority of reflected flux is due to non-local subsurface effects, but as the scale increases, more and more of this flux is deemed to be local.

This scale dependence exists whenever the reflectance of a material is decomposed into local and non-local components (i.e., according to (2.3)), and as a result, such an appearance model is only guaranteed to be accurate for the scale at which it is acquired. Efficient acquisition of multi-scale appearance models remains an open research problem.

3

Spatially- (and Temporally-) Varying Reflectance Models

This section describes the challenges and emerging techniques related to representing the measured appearance of opaque surfaces with spatially-varying (inhomogeneous) reflectance. Recall from Chapter 1 that the appearance of these surfaces is characterized by the Spatially-Varying Bidirectional Reflectance-Distribution Function (SVBRDF) [142] which gives the amount of light reflected from an object's surface along direction ω_o as a function of the position on the surface (x, y) and the direction of incoming light ω_i :

$$S(x, y, \omega_i, \omega_o), \tag{3.1}$$

where these directions are defined with respect to the local surface normal and tangent vectors. Recall also that this function depends on wavelength λ , although most applications in computer graphics ignore this by considering its values within a tristimulus color space (e.g., RGB). As we will see, the fact that this function is defined over a high-dimensional (6D) domain makes it challenging to represent. Finally, it's important to recall the differences between the SVBRDF and the Bidirectional Texture Function (BTF) [24] as discussed earlier. This chapter reviews representations of both.

We begin with a brief review of state-of-the-art devices for acquiring SVBRDFs and BTFs of real-world objects (the previous chapter covered this topic in detail) and then provide a broad review of representations of these measured datasets in chronological order of when they were published. This review is organized around whether or not a technique employs a parametric function of the local surface reflectance or relies instead on a more general-purpose non-parametric model. We then discuss the important role of basis decomposition algorithms in this area and highlight a particular project, the Inverse Shade Tree framework [94], in order to help illustrate key ideas and open research problems. Finally, we conclude by reviewing prevailing representations of measured spatially- and temporally-varying appearance of physical objects.

3.1 Acquisition

Modern image-based systems for capturing the spatially-varying reflectance of real-world objects use digital cameras and point light sources to record measurements sampled densely over the object surface at multiple view and light directions (Figure 3.1, left). For a properly calibrated rig, in which the spectral and angular sensitivity of the camera and output of the source are known, each pixel value corresponds to the integral of the SVBRDF over the region of the surface A subtended by the pixel, the solid angle of incident directions Ω_L subtended by the

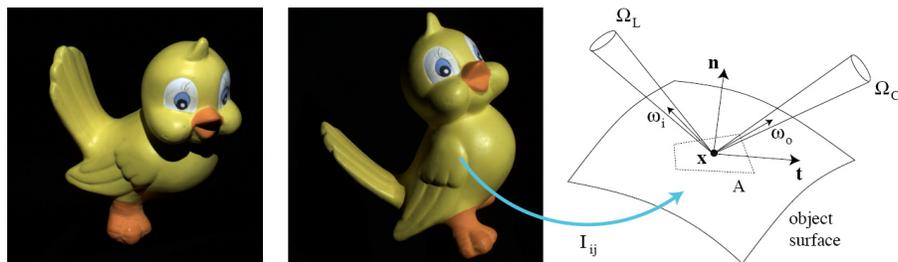


Fig. 3.1 **Left:** Images taken from different viewpoints of an object illuminated by a point light source capture measurements of the SVBRDF defined over its surface. (Images reproduced from [101].) **Right:** The response at each pixel measures the integral of the SVBRDF over the surface area A visible below that pixel, the solid angle Ω_L subtended by the source and the solid angle Ω_C subtended by the camera's lens system.

light source, and the solid angle of reflected directions Ω_C subtended by the camera’s lens system:

$$I_{ij} = \int_{A_j} \int_{\Omega_{L_i}} \int_{\Omega_{C_i}} S(x, y, \omega_i, \omega_o) d\omega_o d\omega_i dA, \quad (3.2)$$

where I_{ij} refers to the j -th pixel in the i -th image. In order to correctly interpret the response at each pixel as a measurement of the SVBRDF two conditions must be met. First, the surface geometry at each pixel must be known — this consists of the average surface position, the surface normal, and, for anisotropic materials, the tangent direction. Additionally, the resolution of the imaging device and the accuracy of this geometric model must result in samples of a physically valid BRDF which conserves energy and obeys Helmholtz reciprocity (Section 2) — inaccuracies in the geometric model is a key source of error. Second, each pixel must measure only light that was directly reflected from the object surface. Light arriving at a surface position which was reflected from another part of the model will contaminate these measurements.

It is useful to categorize acquisition systems based on whether they restrict the target to some simple known geometry (e.g., planar [19, 22, 23, 24, 48, 64, 94, 113, 118, 125, 134]) or allow arbitrary 3D shapes [10, 47, 52, 58, 59, 100, 101, 132, 144, 162, 205, 206]. The former avoids the difficult and error-prone task of acquiring and aligning a model of the 3D geometry of the target object to images like the ones in Figure 3.1, left. On the other hand, most spatially-varying objects do not have a simple geometry nor can they be easily rolled onto, for example, a planar or cylindrical substrate. In these cases, the problem of jointly measuring the shape and surface reflectance is unavoidable. Previously proposed systems capture the target geometry using either optical triangulation-based methods [47, 162, 144] (often after coating the object in a diffuse powder [57, 202]), silhouette-based methods [47, 132], non-optical methods such as computed tomography (CT) [101] or user-assisted photogrammetric techniques [10, 206]. In a second step, the obtained 3D model is aligned to 2D images. Although this registration problem has been heavily scrutinized in the literature [5, 6, 42, 91, 98, 99, 157, 183, 185, 199], it remains challenging in practice. Techniques have also been proposed that recover geometry

and reflectance simultaneously in a single optimization [52, 58, 59]. However, published techniques place restrictions on the reflectance (i.e., that it conforms to a particular analytic BRDF model) and often recover only a partial estimate of the object geometry (i.e., surface normals from a single viewpoint [58, 59] or a height field [52]).

An example of a device suitable for measuring the SVBRDF of physical objects is the Stanford gantry [104]. This system allows placing a camera and light source at arbitrary positions on the sphere surrounding a desktop-sized target object and thus allows sampling all six dimensions of this function.

3.2 Representation

The key benefit of using measured data to render images of the acquired material properties under novel (virtual) lighting conditions and different geometric configurations is that it provides the most accurate model of their visual appearance. However, this accuracy comes at the cost of large storage requirements and processing times. To meet this challenge, researchers have developed a number of representations of the SVBRDF and BTF that reduce the amount of data that must be stored and processed during rendering. Another desirable property of a representation is that it allows editing the angular and spatial components of the measured sample. This is crucial for incorporating measured data into a production setting where designers require control over the materials and objects in a virtual scene.

3.2.1 Parametric Models

A common approach is to model the reflectance at each surface location using any one of the parametric BRDFs reviewed in Section 1. This greatly reduces the amount of data that must be processed since only a handful of values must be maintained at each surface point. Furthermore, this approach is particularly useful when only a small number of input images are available since it provides a way of smoothly interpolating the reflectance at arbitrary view and light directions. The downside of this approach is that it assumes the reflectance conforms to the chosen analytic model which may not always be the case. For

example, the Cook–Torrance BRDF is able to reproduce the appearance of plastics and metals, but it cannot capture other effects such as anisotropy or retro-reflection (backscattering). Furthermore, this method is only appropriate for datasets that provide measurements of the *local* reflectance at each surface point and would not be appropriate for modeling more general BTFs. Another disadvantage, which is not always avoided by alternative approaches reviewed later, is that this approach requires a non-linear optimization to estimate the BRDF parameters that provide a best fit to the measured data. This is often a fragile process that is sensitive to poor local minima [141].

Sato et al. [162] estimate the parameters of the Torrance–Sparrow BRDF over the object surface. They fit the diffuse and specular parameters separately after having separated the input images into these components using a color-space technique. Their final representation allows only smooth spatial variation in the specular fall-off whereas the diffuse color can take on an arbitrary value at each surface point.

The inverse rendering framework proposed by Yu et al. [206] models the spatially-varying reflectance of building interiors using the Ward BRDF. The parameters of this BRDF are estimated from natural images with complex lighting that include secondary (global) illumination effects. They simulate these global effects to isolate measurements of the local surface reflectance. Their final representation allows high-frequency variation in the diffuse component while the specular parameters are estimated at a coarser spatial resolution.

Nishino et al. [144] fit the parameters of a simplified Torrance–Sparrow BRDF to the reflectance measured at a dense set of points over the surface of a 3D object. The shape of the target object is digitized separately using a laser scanner. In addition to the surface reflectance, their optimization also estimates the distribution of incident (distant) lighting in the form of an environment map.

Lensch et al. [100, 101] describe a representation of the spatially-varying reflectance of desktop-sized objects with arbitrary 3D geometry. They apply the Levenberg–Marquardt [103] optimization algorithm along with a divisive clustering technique that assigns each surface point to one from a small set of *component materials*, each represented with the Lafortune BRDF [93]. Their final representation models

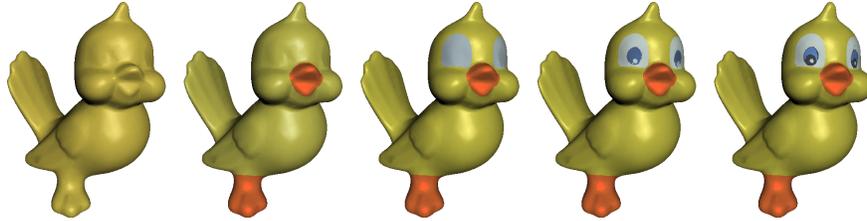


Fig. 3.2 Reflectance clusters computed for the bird model at each step in the divisive clustering algorithm of Lensch et al. [101]. The reflectance of each is represented using the Lafortune BRDF [93].

the reflectance at each surface location as a linear combination of these BRDFs (typically less than five). Note that because this approach effectively shares angular measurements between different surface locations, it avoids potential problems associated with fitting the BRDF parameters independently at surface locations where the reflectance was not adequately sampled (e.g., the specular highlight at a particular surface point may not be observed in any of the input images). Figure 3.2 shows a typical set of cluster centers along with a visualization of their spatial distribution.

Daubert et al. [26] fit the parameters of the Lafortune BRDF to the reflectance *simulated* at a dense set of surface locations over virtual samples of cloth also using the Levenberg–Marquardt algorithm. McAllister et al. [125] and McAllister [126] take a similar approach, also modeling the SVBRDF as a collection of Lafortune lobes. Both projects demonstrated that these fit spatially-varying parameters can be stored as a collection of texture maps and evaluated efficiently within a physically-based or interactive rendering system.

Gardner et al. [48] fit the parameters of the Ward BRDF to a sparse set of angular measurements of roughly planar samples obtained by sweeping a linear light source over the target. Their setup along with the fact that they use planar samples avoids undersampling important features of the BRDF (diffuse amplitude, specular amplitude, and specular fall-off) at any surface location and ensures a stable and robust optimization. They also estimate the degree of translucency over the target surface using a backlight and present applications to document scanning and analysis.

Georghiadis et al. [52] fit the parameters of the Torrance–Sparrow BRDF to the local reflectance at each surface point. In addition to reflectance, they also compute the 3D shape of the target object in the form of a heightfield seen from a single viewpoint. Finally, their optimization treats the positions of the light sources as additional free parameters and thus automates this step in the calibration. However, their final representation requires the sample to have a homogeneous specular lobe (amplitude and fall-off), but does allow arbitrary variation in the diffuse component.

Goldman et al. [58, 59] estimate the surface normals and reflectance in the form of an isotropic Ward BRDF at each pixel in a fixed camera viewpoint from images taken under known variable point lighting. Similar to the method of Lensch et al. [101], they model the per-pixel reflectance as a linear combination of a small number of *basis BRDFs*. This basis, the projection of the measurements into this basis (referred to as *spatial blending weights*), and the surface normals are all estimated together using a standard coordinate descent approach where each subset is alternately optimized using the Levenberg–Marquardt algorithm while the other parameters are held fixed. In addition to greater compression rates and a more stable optimization, this study makes the important observation that expressing the data within a low-dimensional linear basis also allows making intuitive edits to the sample. The parameters of each individual basis BRDF can be adjusted to edit their associated reflectance everywhere they appear over the surface. Conversely, the spatial blending weights can be edited to alter the distribution of these component materials. We return to this important connection between basis decomposition and editability later in Section 3.3.

Marschner et al. [118] model the appearance of polished wood as the weighted sum of a standard specular and diffuse component with an analytic model of the way light is reflected from a dense collection of oriented fibers located below the surface. The fiber orientation is estimated at each surface location through a linear optimization that assumes that the brightest measurements lie along the specular cone and the remaining BRDF parameters are derived using a simple statistical analysis. Light reflected directly from the surface is ignored and

therefore the shape and size of the surface specular lobe must be set manually.

In summary, using an analytic BRDF to model the reflectance at each surface location over a spatially-varying material sample leads to a very compact representation that is easily integrated into existing real-time and physically-based rendering systems. Another important advantage is that this approach can be effective when only a small number of input images are available. The key drawback is the loss of accuracy when the measured reflectance substantially deviates from the chosen analytic model and the often fragile non-linear optimization that must be performed.

3.2.2 Non-Parametric Models

Alternative approaches to representing the SVBRDF and BTF use general-purpose dimensionality reduction and function interpolation algorithms. These so-called *non-parametric* methods do not require that the measured reflectance comply with a particular analytic function and, as a result, apply to a wider range of materials. Another benefit of these methods is that the underlying optimization is often more stable and less susceptible to poor local minima. Their increased generality and accuracy, however, often come at the cost of requiring a denser sampling of the angular domain and can run counter to the goal of obtaining an intuitive representation that lends itself to direct editing.

Common to these methods is the fact that they all exploit phenomena present in real-world reflectance functions. Indeed, requiring that the measured data conform to a particular analytic function is one example of this philosophy, but there are less restrictive constraints that allow achieving a practical representation while also applying to a wide class of materials. We have already discussed *isotropy* and *reciprocity* as being common, and useful, properties of reflectance from an acquisition standpoint because they substantially reduce the angular domain and thus the number of required images. These traits are exploited to achieve compact representations as well. Real-world SVBRDFs also tend to be *compressible* in the sense that they permit

a sparse representation when projected into particular bases such as the wavelet basis [123]. However, the degree of compressibility is often heavily influenced by the way in which the angular domain is parameterized as discussed in Chapter 1. Another common property is *separability*, which refers to the fact that spatially-varying reflectance can often be written as a linear combination of diffuse and specular components. This is useful because when isolated, each of these components exhibits different (and exploitable) behavior. For example, the diffuse component is typically well-represented by a Lambertian model and can often be reliably estimated over the entire surface from just a handful of images. *Spatial smoothness* is another important property which refers to the fact that for many surfaces, reflectance is slowly varying from point to point. This is more often true for the specular reflectance component (e.g., [162, 212]). Thus, knowledge of the reflectance at one point on a surface is often a good indicator of the reflectance at another. Finally, *spatial regularity* is another way of describing the correlation between the reflectance at distinct surface points on the same surface. Here, it is assumed that the reflectance at all surface points can be written as linear combinations of a single set of basis BRDFs (e.g., [59, 94, 100]).

Since the work of Dana et al. [24], Nishino et al. [143], Cula and Dana [20], Zalensny and Van Gool [207], and Furukawa et al. [47], standard rank-reduction algorithms such as Principal Component Analysis (PCA) have been used extensively to compress measured SVBRDFs/BTFs of physical samples [89, 109, 135]. This requires first resampling the data along a regular grid of angular positions and organizing these samples into a matrix that is subsequently factored using the Singular Value Decomposition (SVD) algorithm. Local PCA has also been used to compress BTFs [133]. This involves projecting subsets of the data into different linear bases designed to achieve a greater compression ratio. Vasilescu and Terzopoulos [184] and Wang et al. [187] extended this work to consider multilinear representations that are computed by applying the N -mode SVD algorithm to higher-order tensors. This is a more natural approach since these functions are inherently high-dimensional and allow exposing (and exploiting) coherency between all of the dimensions at once. Leung and Malik [102] used the

k -means clustering algorithm to compress BTFs. In their approach, a dataset is represented as a small collection of *textons*, which capture the spatial and directional appearance of small surface patches. Weistroffer et al. [192] avoid resampling the input data by computing the factorization with respect to a secondary basis defined over the 4D BRDF domain. They show examples of using radial basis functions and a database of measured isotropic BRDFs [122] as the secondary basis to model the samples captured by Lensch et al. [100].

Malzbender et al. [116] introduced Polynomial Texture Maps (PTMs) which represent the directionally dependent appearance of an object at each pixel in a fixed camera viewpoint as a low-dimensional polynomial. They demonstrated a number of applications of this model that included BTF samples in the sense that the underlying geometry is only approximately known.

Zickler et al. [212] introduced *reflectance sharing* as a general strategy for interpolating sparse measurements of the SVBRDF by fitting a set of radial basis functions. The input is a set of images of a known three-dimensional shape, with each image being captured under collimated illumination in a known direction. To exploit spatial smoothness, they view each pixel as a sample lying in the (5D isotropic) SVBRDF domain, and note that each image provides a near-continuous sampling of a 2D slice in this domain. SVBRDF estimation is formulated as a scatter-data interpolation problem, in which samples are simultaneously interpolated in both the angular and spatial dimensions.

The Inverse Shade Tree framework introduced by Lawrence et al. [94] is another non-parametric approach for modeling spatially-varying reflectance. This project builds on earlier work [86, 95, 127] that demonstrated measured (tabulated) BRDFs can be factored into a small number of low-dimensional components using algorithms like SVD. They focus on the goal of intuitive editing and this work is discussed in greater detail later.

Wang et al. [190] model the reflectance at each surface point of a spatially-varying surface as the product of a non-parametric (tabulated) half-angle distribution and analytic shadowing/masking and Fresnel components. These types of “hybrid” BRDFs were introduced by Ashikhmin et al. [3] and later refined by Ngan et al. [141] where

they were shown to be successful in representing real-world anisotropic surfaces (as compared to available parametric models). The half-angle distribution and tangent direction at each surface point are estimated incrementally using a technique akin to texture synthesis algorithms.

Alldrin et al. [2] present a photometric technique that estimates the surface normal below each pixel in a fixed camera viewpoint along with the surface reflectance represented as a linear combination of a small number of tabulated bivariate BRDFs. Prior work [158, 172] has demonstrated that real-world isotropic BRDFs can be accurately approximated as functions over a properly chosen 2D domain. Alldrin et al. build on this observation by modeling the basis BRDFs as sampled (non-parametric) functions defined over these bivariate domains. This allows achieving greater flexibility and generality. They employ a factorization method similar to that used in the Inverse Shade Tree framework of Lawrence et al. [94].

The following section examines the increasingly important role basis decomposition has come to play in representing measured SVBRDFs. We highlight the Inverse Shade Tree framework [94] as a way to illustrate the key ideas in this area of work and indicate open research questions.

3.3 Case Study: Basis Decomposition and the Inverse Shade Tree Framework

Although (3.1) allows a unique BRDF at every surface location, real-world surfaces are typically composed of only a handful of *component* or *basis* materials distributed in potentially complex spatial patterns. For example, the bird figurine in Figure 3.1 was made from only a few different types of paint. This property may be exploited by projecting the measurements into a low-dimensional linear basis spanned by K functions $\rho_k(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ (which are defined over the BRDF domain) for which $T_k(x, y)$ encode the coordinates of the measurements within this basis:

$$S(x, y, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \approx \sum_{k=1}^K T_k(x, y) \rho_k(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o). \quad (3.3)$$



Fig. 3.3 Images of a greeting card taken from different viewing angles and under varying point illumination.

By way of example, consider the greeting card in Figure 3.3 which was measured at a spatial resolution of 470×510 and angular resolution of 5×400 camera \times light positions using a spherical gantry [104] for a total of 2,000 images (5.5 GB). This dataset is composed of only three distinct materials: two types of colored paper (yellow and blue) and a metallic paint that takes the shape of a dove. The IST framework exploits this type of structure by computing a hierarchical decomposition of these input images into a tree-structured collection of lower-dimensional functions that correspond to intuitive (latent) features in the dataset. Equation 3.3 corresponds to the top-level decomposition that rewrites the 6D input SVBRDF as the sum of products of a set of 4D functions (basis BRDFs) and 2D spatial blending weight maps (coordinates in this basis). As illustrated in Figure 3.4, the complete process, which involves further decomposing each of these tabulated 4D basis BRDFs into 2D and finally 1D functions, can be thought of as computing a *shade tree* [17] whose leaf nodes capture elements such as the shape and size of the specular lobe and the spatial distribution of each component material over the object surface. This process accomplishes two things. First, it compresses the input because only a small number of lower-dimensional (1- and 2D) functions must be stored instead of the densely sampled 6D input. Second, if the factorization at each level produces a meaningful and physically valid decomposition then the final representation will allow making intuitive edits

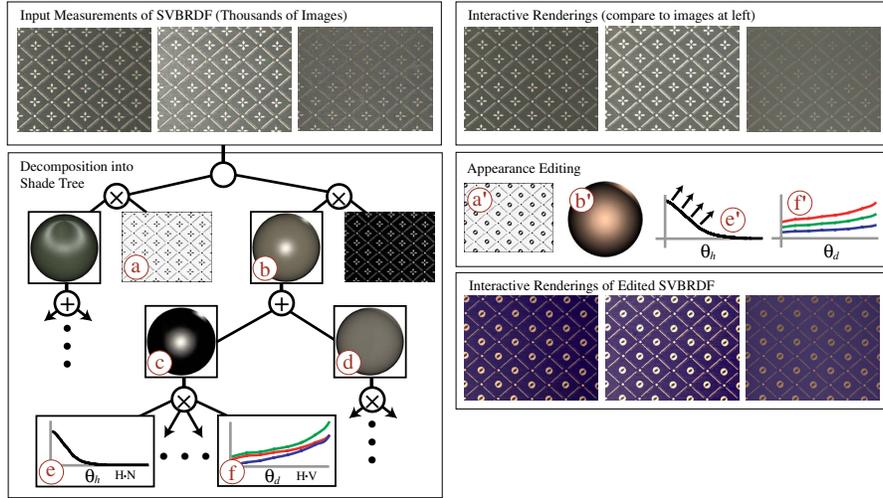


Fig. 3.4 The Inverse Shade Tree framework [94] is a set of techniques for decomposing a measured SVBRDF into a set of (a) spatially-varying blending weight maps and (b) basis BRDFs. The basis BRDFs are factored into sampled 2D functions corresponding to (c) specular and (d) diffuse components of reflectance (we show lit spheres rendered with these factors, not the 2D factors themselves). These 2D functions are further decomposed into (e & f) 1D curves. In addition to providing accurate interactive rendering of the original SVBRDF, this representation also supports editing either (a') the spatial distribution of the component materials or (b') individual material properties. The latter is accomplished by editing (e' & f') the sampled curves. (Figure reproduced from [94].)

to the angular and spatial components of the measured input. The key research challenge is developing an algorithm that can reliably compute such intuitive decompositions. It's illustrative to compare alternative factorization methods with this goal in mind. We do so below, focusing on the top-level decomposition expressed in (3.3) where the goal is separating the measured input into a set of intuitive and physically valid component BRDFs and their respective spatial distributions.

PCA: Perhaps the most popular rank-reduction algorithm is PCA, including extensions such as multilinear tensor factorization [184]. The main advantage of PCA is that it yields a *global* minimum in the sense of total least squares. However, this algorithm recovers a basis that is orthonormal and rarely provides a *meaningful* description of the input. In particular, this algorithm produces negative values which are not present in a physically valid BRDF.

Clustering: A popular method for clustering data is the k -means algorithm [68]. Like all clustering algorithms, k -means partitions the input into disjoint sets, associating each point (row in the aforementioned data matrix) with a representative *cluster center*. This set of cluster centers can be interpreted as spanning a linear subspace into which the data is projected as with the approach of Lensch et al. [100]. Although clustering performs well on datasets for which the component materials are well-separated over the surface, it struggles when these materials are only ever observed in combination with one another such as, for example, datasets that include semi-transparent materials.

Non-Negative Matrix Factorization: Another matrix decomposition approach is Non-Negative Matrix Factorization (NMF) [96]. As with related algorithms such as Probabilistic Latent Semantic Indexing [76], NMF guarantees that both resulting factors contain only non-negative values. One motivation for this constraint is to encourage the algorithm to describe the input data as the sum of positive *parts*, thereby producing a more meaningful factorization. However, this non-negativity constraint alone is often not enough to achieve an intuitive separation.

Alternating Constrained Least Squares (ACLS): Lawrence et al. [94] introduce ACLS, a straightforward coordinate descent technique that allows computing a matrix factorization subject to general linear constraints. Building on the work of Goldman et al. [58, 59] they also observe that even when a sample consists of multiple component materials only a few typically contribute to the reflectance at any one surface point. This notion of sparsity is enforced by modifying the objective function to favor blending weights that are closely aligned to one of the coordinate axes in the computed basis.

Figure 3.5 visualizes the set of blending weights produced by these factorization algorithms when applied to a dataset composed of a plank of oak wood with a semi-transparent piece of tape and a bicycle retroreflector. The visualization of PCA displays positive values as green and negative values as red; the other methods produce non-negative weights which are shown as grayscale images. Note that although PCA gives an optimal result in terms of this error metric, it fails to provide an intuitive representation. For example, consider the difficulty

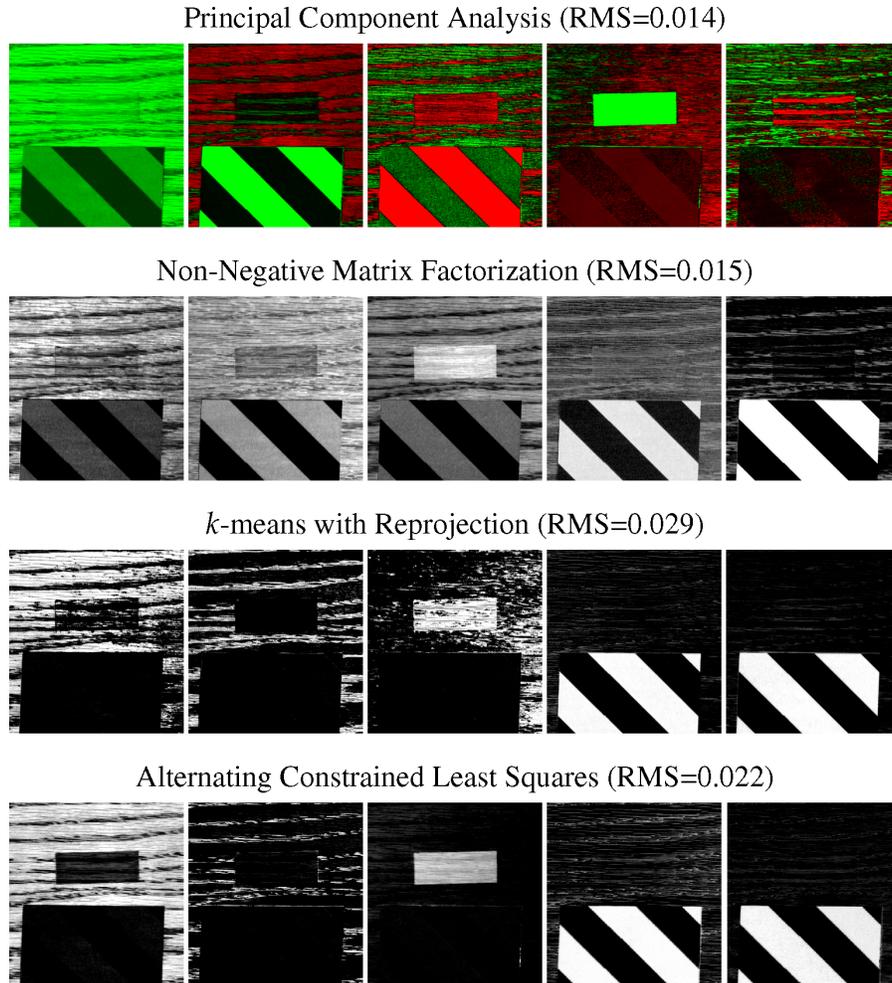


Fig. 3.5 Visual comparison of the spatial blending weights computed using different factorization algorithms for the “Wood+Tape” dataset. ACLS (bottom row) enforces sparsity and guarantees that the component BRDFs are physically valid (energy conserving and reciprocal). This aids in providing automatic separation of the measured data into its correct component materials.

in using this representation to change only the properties of the light wood grain. NMF [96] produces a more understandable separation, but it also displays a significant amount of undesirable blending between the component materials which would complicate editing. Clustering

algorithms provide a representation that is as sparse as possible in the way defined previously, but struggles to separate a clear term for the semi-transparent tape since this is only ever observed in combination with the underlying wood. In this example, ACLS is able to more cleanly isolate the component materials, but at a slightly larger numerical error.

In summary, the goals of compressing measured SVBRDFs and exposing their intuitive latent structure are complimentary. The first is necessary to allow efficient rendering and the second enables a designer to edit and refine the measured sample. Most representations that achieve both of these goals do so by projecting the input into a suitably constructed basis through some form of constrained optimization. However, the AppWand system introduced by Pellacini and Lawrence [151] describes an alternative approach. In this system, the user indicates regions of the dataset they wish to modify using a stroke-based interface. These sparse constraints are propagated throughout the entire dataset using an interactive optimization that enforces the policy that regions with similar appearance receive similar edits. They show results on datasets modeled as individual parametric BRDFs at each surface location. Investigating the advantages and disadvantages between decomposition-based and propagation-based editing techniques and bringing them to a wider range of material functions is an active area of research.

3.4 Space- and Time-Varying BRDFs (STVBRDFs)

A number of recent research projects have focused on capturing and representing the space- and time-varying reflectance of real-world surfaces. Consider, for example, the complex spatial and temporal patterns present in the appearance of a banana peel as it decays or for a metallic object as its surface oxidizes.

Gu et al. [62] used a “light stage” setup [30, 33, 84, 193, 195] to measure the STVBRDF of physical samples. These devices consist of a set of multiple lights and cameras rigidly mounted to a geodesic dome that can capture a single multi-light/multi-view dataset in several seconds. They measured the spatially-varying reflectance of each

sample at different points in time, for example, as copper rusted, toast burned, or a banana rotted. They fit the Cook–Torrance BRDF to the measurements at each spatio-temporal location and analyzed how these parameters changed over time at each spatial location. Based on this analysis, they compute a separate temporal curve for each BRDF parameter (diffuse color, specular color, specular fall-off) under the assumption that all of the measurements show the same physical process at potentially different points in time. This final representation allows re-targeting the measured appearance on arbitrary objects and editing the space–time trends to produce custom patterns.

Wang et al. [188] introduced *appearance manifolds* to model the time-variant appearance of real-world materials. They used a setup similar to Gu et al. and fit the parameters of the Ward BRDF to the data at each space–time sample location. They treat each vector of fit parameters as a point in a high-dimensional space and hypothesize that the collection of vectors within a single dataset form a 1D non-linear manifold (curve). They use a method similar to IsoMap [177] to reconstruct the best fitting 1D curve through these vectors and apply a texture synthesis algorithm to generate coordinates along this curve that allows applying the measured reflectance in novel spatial/temporal patterns over arbitrary geometry.

Context-aware textures were introduced by Lu et al. [110] to model the correlation between material properties and the underlying geometry. For example, the degree of weathering or oxidation at a surface location is often influenced by the local geometric properties such as curvature and accessibility to the surrounding environment. Computing the relationship between the reflectance and the local geometric context allows re-targeting the measured data to objects with arbitrary geometry with user-defined aging factors. They measured the space- and time-varying diffuse albedo of several objects using a 3D scanner and digital camera. They focused on time-varying processes that could be periodically arrested (e.g., they used a chemical reagent to control the oxidation process of copper) so that the spatially-varying reflectance could be measured at fixed points in time.

4

From BSSRDF to 8D Reflectance Fields

In this survey, we have so far concentrated on capturing surface reflectance where incident light is scattered locally at the point of incidence. The models that have been presented ignore global effects such as subsurface scattering, transmission, or interreflections in complicated surface geometry.

4.1 Capturing BTFs

Some of these effects can be captured by a BTF that records the apparent SVBRDF of a surface sample patch. In contrast to an SVBRDF the bidirectional texture function (BTF) is used to represent the effects of self-shadowing and self-occlusions which are common to surfaces with complicated geometric structure such as woven fabrics.

To acquire the spatially varying reflection behavior, one needs to capture one image of the given surface patch for every pair of incident and outgoing directions, i.e., from each point on the hemisphere a set of pictures has to be acquired where the light sources cover all possible incident directions.

The acquisition of BTFs from simulated data has been demonstrated in [180]. Most commonly, moving robot gantries are employed to

cover the multitude of camera and/or illumination directions with few devices [23, 30]. Instead of moving cameras and light sources, one can keep the camera fixed and tilt the surface for addressing the individual directions [134]. Of course, the hemisphere of direction can further be populated by a set of light sources at fixed locations [37, 132, 193] speeding up the capture process as switching between light sources is faster than moving a light source. Similarly, a set of cameras can be used to capture the necessary data in parallel for all viewpoints [132]. This comes at the expense of setting up, calibrating, and maintaining multiple cameras.

In order to cover a larger portion of the hemisphere with a single device, setups involving mirrors and projectors have been proposed for capturing BTFs, e.g., [25, 54, 64]. Han and Perlin [64] arranged three mirrors as a kaleidoscope and set up a single camera and a single projector via a beam splitter to share the same view. Looking into the kaleidoscope with a camera the image of the surface patch appears multiple times in it, each sub-image showing the surface from a different viewpoint. Due to the multiple reflections, one has to account for attenuation. Similarly, the illumination from a projector can be tuned to illuminate the surface from different angles by selecting the region of the corresponding sub-image in the projector. This way, the 6D BTF or reflectance field can be acquired efficiently.

As the raw data of a BTF acquisition is quite large for a single surface patch a number of compression schemes based on spherical harmonics, local and global principal component analysis (PCA), tensor decomposition and others have been proposed (see previous chapter and [134]). The choice of representation will further influence the following rendering algorithm which in the meantime can be evaluated at real-time rates for environmental illumination [129, 134, 163].

4.2 Acquiring BSSRDFs

BRDFs and BTFs treated so far assume that the light hitting the surface is parallel, originating from an infinitely far away light source. If one wants to simulate a close by light source, or, equivalently, to project a light pattern into the scene, e.g., from a spot light or as the result of

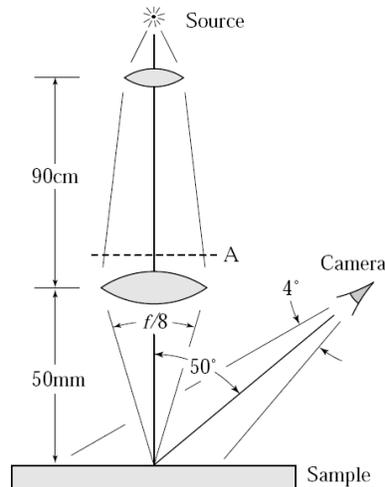


Fig. 4.1 Acquisition setup of Jensen et al. [83] for measuring homogeneous BSSRDFs.

two objects interacting with each other, the far-field assumption is not valid. Instead, one needs to record the so-called near-field reflectance field which couples incident to outgoing light *rays* rather than directions. Here, the reflectance depends both on the point of incidence and the reflection point.

This reflection behavior is captured by a BSSRDF or an 8D reflectance field. For a dense sampling of this 8D function one would need to record the reflected 4D light field for every possible incident light ray (4D), i.e., one needs to sample how any incident light ray might influence any outgoing light ray even though they might not intersect.

Because of their high dimensionality BSSRDFs are in general hard to represent and to acquire. A simple analytic BSSRDF model for homogeneous materials has therefore been proposed by Jensen et al. [83]. It has later been updated to incorporate homogeneous multi-layer materials [34]. In the first paper, a simple measurement setup is presented for estimating the parameters of the model for a single homogeneous material sample: A focused light beam hits the homogeneous slab at one point, and the resulting spatial distribution of the reflected light is measured using a camera. A simpler device for instant measurement of a few samples has been proposed in [195].

4.3 Diffuse Subsurface Scattering

For heterogeneous surfaces it is impractical to densely sample all eight dimensions of the BSSRDF/reflectance fields. In order to allow for reasonable sampling effort, one strategy is to assume a less complex light transport.

Goesele et al. [57] present a measurement setup for acquiring the appearance of translucent objects with a high scattering albedo. In these cases, a photon traveling some distance through the material undergoes so many scattering events that the incident light direction has actually no influence on the outgoing light direction. Since the directional dependence can be dropped from the full 8D BSSRDF, the problem can be reduced to a 4D diffuse scattering $R(x_i, y_i, x_o, y_o)$ function that solely depends on the point where the light enters the material and the position where it leaves when being reflected. For any pair of points R indicates how much the incident irradiance at point (x_i, y_i) contributes to the outgoing radiosity at point (x_o, y_o) .

A simple 4D tensor can be used to represent this 4D function. In order to measure its entries Goesele et al. make use of a laser projector, that sweeps an individual light point over the surface of a translucent object (see Figure 4.2). A set of high-dynamic-range (HDR) video cameras captures the reflected light at every other surface point. One of these measurements corresponds to exactly one slice of the 4D tensor. Illuminating every surface point once eventually fills the entire tensor.

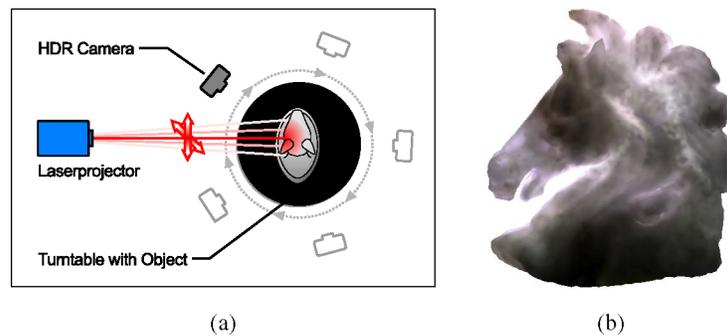


Fig. 4.2 Acquisition setup (a) of Goesele et al. [57] for measuring the appearance of heterogeneous translucent objects such as this alabaster figurine on the right (b).

Because of occlusions and self-shadowing, it is, however, likely that for some parts of the object no measurements are available, in which case texture inpainting is applied to fill the gaps. For efficient storage, a hierarchical representation of the tensor is chosen, providing high sampling rate only close to the point of incidence, where the BSSRDF drops off quickly, while for distant points a coarser sampling is sufficient.

Recently, Peers et al. [150] presented a different setup to acquire the spatially varying 4D diffuse BSSRDF of a planar slab of material. In order to accelerate the acquisition, a grid of points is swept over the surface. A PCA of the acquired tensor allows for a high compression and to extract the principle behavior of the BSSRDF such that it can be transferred to different object geometry. When doing so, an important step is to rearrange the tensor in such a way that the point of incidence for any row is shifted to exactly the same column. Although the heterogeneity will influence the drop of intensity with distance to the point of incidence in a different way for each surface location, the principle behavior, roughly following an exponential function, is nicely aligned in this way allowing for higher compression rates.

Heterogeneous BSSRDFs have been modeled and rendered efficiently by Wang et al. [189] using the diffusion equation. Here, the diffuse light transport through the medium is explicitly modeled in the volume. With acquisition patterns similar to [150], they acquire either the transmission or the back-scattered component of the subsurface scattering and solve the inverse diffusion problem, initialized by assuming a homogeneous medium. To accelerate the evaluation in the forward and backward steps a polygrid diffusion algorithm is proposed.

4.4 Arbitrary Light Transport

For arbitrary materials and scenes, Masselus et al. [121] presented the first acquisition system for reflectance fields that is suitable for relighting with 4D incident light fields, i.e., where the reflected light depends on individual incident light rays. For the acquisition, a video projector sweeps a small block of light over the scene. In order to cope with the complexity of the acquisition problem the appearance is captured for a single viewpoint only. Additionally, the resolution of the incident light

field is limited to a projector resolution of only 16×16 for a couple of projector locations. This low resolution results in clear block artifacts in the relit images.

4.4.1 Single View — Single Projector

In order to avoid those artifacts, it is necessary to measure the reflectance for every pair of rays between a camera and a projector, i.e., to acquire the reflectance for every pair of camera and projector pixels, again resulting in a fourth-order tensor. While, in principle, this high-resolution reflectance field could be acquired using scanning, it would be a rather slow process.

Sen et al. [166] exploited the fact that for quite a number of real-world scenes the light transport matrix/reflectance field is rather sparse, i.e., that only a small fraction of the possible input rays actually contribute to the same reflected ray. In this case, it is possible to exploit the sparseness by illuminating the scene and measuring the reflected light rays for multiple illuminating light rays at once. It is possible to turn on two light rays/two projector pixels at the same time and tell their corresponding measurements apart when these two rays affect completely separated parts of the scene/the camera images. Let us call such two illumination rays **radiometrically independent**. In the same way one can also call two blocks of the projector pixels radiometrically independent if no camera pixel will be illuminated by both at the same time. Because of this property it is possible to measure the rays inside the independent blocks in parallel, i.e., to parallelize the exact acquisition of these two blocks. In their study, Sen et al. propose a hierarchical approach for determining which sub-blocks are independent: Starting from a full white projector image, each block is subdivided into four children which again get subdivided. Initially, this will require one measurement per block corresponding to a sequential acquisition. At some point in time the algorithm might, however, detect that at some level two blocks are now radiometrically independent, allowing for parallelized subdivision of these blocks in the future. The net effect of this parallelization is significant, resulting in an observed complexity that is $O(\log(n))$ for n projector pixels. For quite a number

of scenes the pixel-to-pixel reflectance field between a one mega-pixel projector and a camera can be acquired in only a couple of thousand images instead of a million.

Once having acquired the pixel-to-pixel light transport one can apply **Helmholtz reciprocity** to invert the role of projectors and cameras. Helmholtz reciprocity states that the reflectance measured for one path does not change no matter if one follows the path from the light source to the receiver or the other way around. One simply has to compute the transpose of the acquired tensor to obtain the reflectance field from the camera (the new virtual projector) to the projector (which gets the new camera). The transpose corresponds to just a resorting of rays, and therefore can be computed very efficiently (see Figure 4.3). The process is called **dual photography**.

This dual imaging paradigm can be used to efficiently capture a 6D reflectance field from a single viewpoint, i.e., to measure the projector to camera reflectance fields for multiple projectors. The problem is that during the acquisition the reflectance fields have to be captured

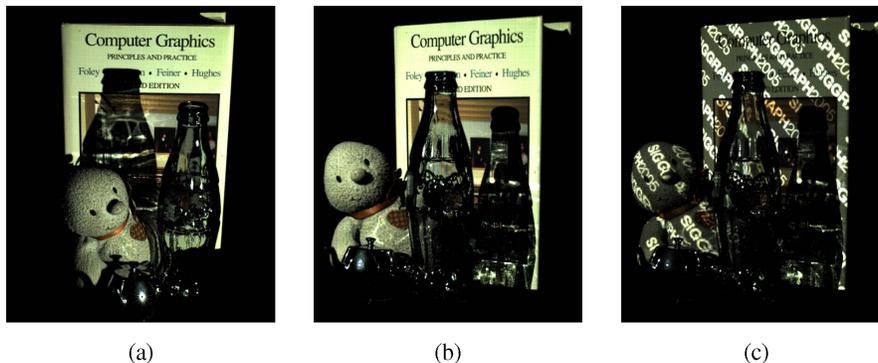


Fig. 4.3 Dual photography: (a) conventional photograph of a scene, illuminated by a projector with all its pixels turned on. (b) After measuring the light transport between the projector and the camera using structured illumination, dual photography is able to synthesize a photo-realistic image from the point of view of the projector. This image has the resolution of the projector and is illuminated by a light source at the position of the camera. The technique can capture subtle illumination effects such as caustics and self-shadowing. Note, for example, how the glass bottle in the primal image (a) appears as the caustic in the dual image (b) and vice-versa. Because we have determined the complete light transport between the projector and camera, it is easy to relight the dual image using a synthetic light source (c).

sequentially for each projector because projectors are active devices. Their projected patterns might actually illuminate the same points in the scene causing difficulties when trying to separate their contribution. If one uses the dual setup, however, where the original camera is replaced by a single projector and all projectors are replaced by cameras, one can very well acquire the projector/camera reflectance fields in parallel since cameras are passive devices which do not interfere with each other [166]. Applying Helmholtz reciprocity, this setup can virtually be transformed into the single camera/multiple projector configuration. By swapping camera and projectors one can capture a 6D reflectance field at the same time cost as a 4D reflectance field. The resulting data now allow to relight an arbitrary complex scene with arbitrary incident light fields, i.e., with high-frequency illumination patterns from various virtual projector positions.

Following the same idea of exploiting Helmholtz reciprocity, Hawkins et al. [69] propose a dual light stage, where a high-power laser projector is illuminating the scene from one viewpoint (turned into the virtual camera location later), while a camera with a fish-eye lens is observing the reflections captured in a diffuse sphere surrounding the object. As the diffuse sphere allows for a dense sampling of the outgoing hemisphere (incoming after applying reciprocity), the system is very well suited to capture the complex reflections and caustics due to transparent or mirroring objects due to illumination from an environment map.

4.4.2 8D Reflectance Fields

While methods for measuring 6D reflectance fields such as [69, 121, 166], can in principle be extended to measure the 6D reflectance fields from a set of different viewpoints, thereby capturing an 8D field, they inherently would require an acquisition time linear to the number of viewpoints. Garg et al. [49] have proposed the first method which is inherently acquiring an 8D reflectance tensor for general scenes. It can deal with dense tensors and parallelizes the acquisition even between viewpoints. It exploits a symmetric measurement setup and \mathcal{H} -matrices to exploit the coherence in the reflectance data already during the capturing process.

The previously described acceleration for measuring the light transport between a projector and a camera based on the parallelization of radiometrically independent blocks is unfortunately limited to scenes where the light transport tensor is sparse. This is often the case for an individual object in a black room where few interreflections and little subsurface scattering take place. For more general cases, the light transport matrix is typically dense, i.e., every projector pixel indirectly affects every camera pixel due to multiple scattering events. The resulting light transport is, however, rather smooth for large blocks of the tensor. For example, illuminating two neighboring spots on a wall will have a similar effect to all points on an opposite wall. While this smoothness might be partially destroyed by textures on both walls, the underlying light transport still has rather low complexity or low dimensionality — it is called data-sparse. In other parts in the ray space, however, for example, for direct reflections or refractions, the reflectance field might not be smooth at all.

\mathcal{H} -matrices [63] are an efficient way for representing tensors which are partially data-sparse. In an \mathcal{H} -matrix the original matrix is hierarchically subdivided into smaller blocks, e.g., using a quad-tree for a 2D matrix, and for every sub-block a low-rank approximation is given, approximating the original matrix's entries. If the approximation error for one block is too large, the block is further subdivided. As \mathcal{H} -matrices have been originally developed to solve integral equations more efficiently, and since the **Rendering Equation** which describes the light transport in arbitrary scenes is an integral equation, reflectance fields can be very efficiently described by this data structure.

Besides resulting in a compact representation of a reflectance field, \mathcal{H} -matrices can efficiently be evaluated during relighting, where the incident light field is simply multiplied with the tensor.

\mathcal{H} -matrices further open the way for efficient acquisition of reflectance fields of arbitrarily complex scenes where interreflections and scattering cannot be neglected, as well as for the acquisition of 8D reflectance fields.

Garg et al. [49] have proposed a measurement setup that forces the captured reflectance tensor to be symmetric. In their setup every camera is paired with one projector using a beam splitter in such a way

that it is possible to emit light and to measure light exactly along the same ray. In the resulting transport tensor every off-diagonal sub-block is therefore represented twice, once in its original form and once being transposed, i.e., we could capture the original and the dual image for one sub-block with just two images by fully illuminating the corresponding two projector blocks.

Since one of the images corresponds to the sum along the rows of the block and the other image to the sum along the columns of the block, it is possible to obtain a rank-1 approximation of this block with just these two images, simply as the tensor product of the two measurements obtained when first illuminating with one block of one projector, measuring the result in one block of some camera, and then measuring the transpose, i.e., measuring at the block of the first projector and emitting light from the first camera's block.

Let's look at a very simple example where the off-diagonal block B_2 has been determined to be rank-1:

$$T = \begin{pmatrix} B_1 & B_2 \\ B_2^T & B_3 \end{pmatrix} = \begin{pmatrix} 0 & B_2 \\ B_2^T & 0 \end{pmatrix} + \begin{pmatrix} B_1 & 0 \\ 0 & B_3 \end{pmatrix}. \quad (4.1)$$

In this case, we can determine all entries in B_2 (and B_2^T) from just two images while B_1 and B_3 might require additional investigation.

From the intended solution one can subtract the already determined matrix $\begin{pmatrix} 0 & B_2 \\ B_2^T & 0 \end{pmatrix}$, which leaves us with some residual which only contains the remaining, yet unknown blocks. These two blocks are, however, arranged in a very interesting configuration: they are radiometrically independent, since they clearly affect completely different camera and projector regions. As a consequence, these two blocks can again be investigated further in parallel. This **symmetric photography** method allows for the efficient and parallelized acquisition of even dense matrices as long as the matrices are data-sparse.

In Figure 4.4, we show a low-resolution 8D reflectance field for 3×3 cameras and 3×3 projectors. With the symmetric photography approach one can acquire the light transport of scenes as complicated as this glass of gummy bears and faithfully reproduce the appearance of the original object (Figure 4.5).

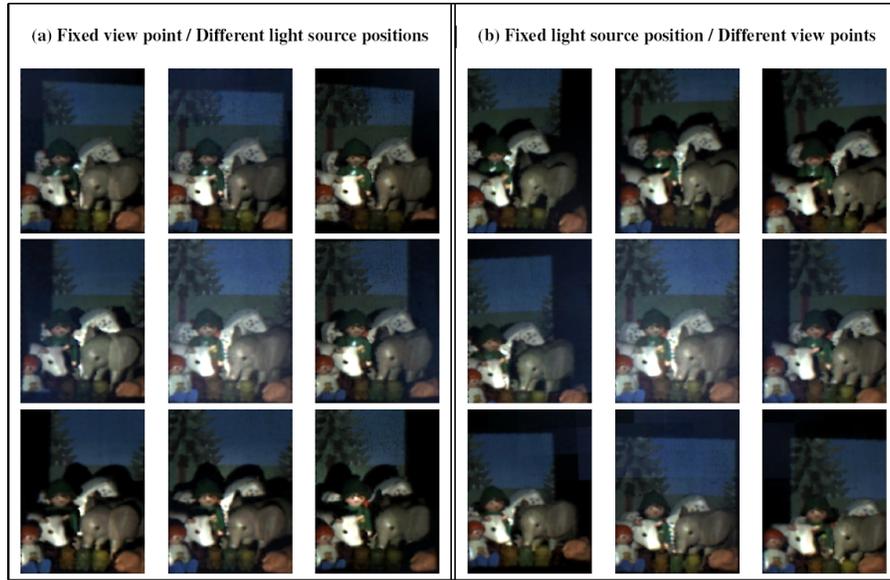


Fig. 4.4 An 8D reflectance field acquired using symmetric photography. The scene has been recorded from 3×3 different viewpoints and can be relit from 3×3 different projectors with full resolution.

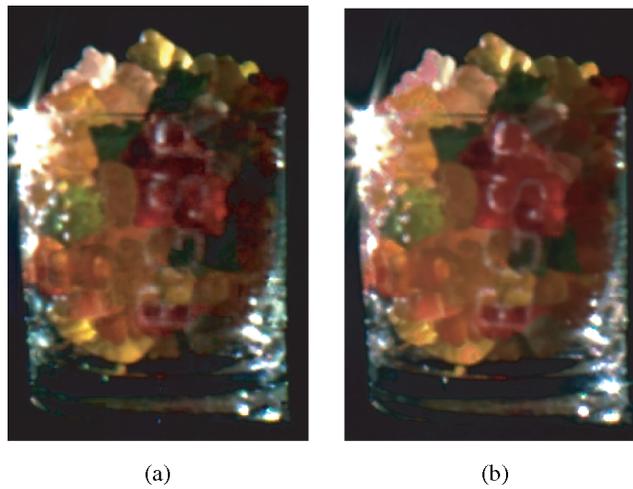


Fig. 4.5 Using symmetric photography the light transport in even very complex scenes can be efficiently captured. The synthetically relit reflectance field (a) matches the appearance of the original object under the same pattern (b).

4.5 Time-Resolved Capture and Interpolation

Besides capturing the full 8D reflectance field of an object, other extensions of the 6D capture process for example focus on capturing time-varying reflectance fields.

The main challenge here is to acquire the 6D reflectance field for each time step. Exploiting specially designed illumination patterns, Hawkins et al. [70] reduced the time to acquire a static reflectance field to a few seconds. The authors demonstrated how individual facial expressions can be acquired in this setting, and how relightable facial animations can be rendered by blending between the captured images.

Even faster acquisition rates have been achieved by employing high-speed cameras. Wenger et al. [193] demonstrated how the time-varying reflectance field of a multi-second performance can be captured. Although at each time step only a single illumination setting can be present, they successfully demonstrated how optical flow between the individual frames can be used to interpolate the reflectance values for the missing illumination directions from neighboring frames [37].

One step further, Peers et al. [149] even estimated the correspondence between two different individuals to transfer the reflectance field captured for one actor onto the current facial expression of another actor. In order to establish the optical flow, both persons are captured in a more or less homogeneous illumination. Then, the intended illumination is used to relight the actor for which a static reflectance field has been acquired. This data is then warped according to the optical flow and then combined with the currently frame of the second actor using quotient images [168].

Optical flow in the illumination domain has further been exploited to increase the sampling density in reflectance fields [14, 45]. In [14], the information in a 6D reflectance field between projectors and a camera is used to estimate the pixel-to-pixel correspondence between the projectors of two neighboring positions. This way, it is possible to warp the illuminating pattern for any virtual in-between projector location in such a way that a smooth motion between the two projectors can be simulated. Fuchs et al. [45] proposed different interpolation schemes to interpolate the point light source positions in 4D reflectance fields,

which most often are acquired by sparsely sampling the hemisphere of illumination directions [44]. While diffuse and smooth reflectance effects can simply be blended, the appearance of specular highlights is interpolated using optical flow, i.e., interpolating the direction of the reflection vector in the environment map. The interpolation scheme allows for example for continuous reflections in mirroring objects or continuous refractions in glass although only a sparse sampling of light directions is provided.

4.6 Conclusions and Future Work

In this chapter we have introduced the notion of reflectance fields for relighting with spatially varying illumination patterns: from the acquisition of heterogeneous translucent objects to methods for acquiring the ray-to-ray light transport in arbitrary materials and scenes.

4.6.1 Open Problems

One big problem of sampling BSSRDFs or reflectance fields is the limited resolution with respect to the incident and outgoing directions. While solutions have been proposed to increase the resolution of the incident illumination [69, 44] by using special light source arrangements, the resolution of the viewing directions is still limited to the spacing between adjacent camera positions. A scheme for adaptively controlling the resolution in the viewing and the illumination direction still needs to be invented.

For representing reflectance fields, various bases have been proposed, wavelets, spherical harmonics, or the above-mentioned \mathcal{H} -matrices. It remains to be seen how to select the optimal representation and how to determine the dimensionality of the light transport locally.

A still outstanding goal is the acquisition of reflectance fields for relighting with 4D incident light fields for dynamic objects. While initial solutions to measuring time-varying far-field reflectance fields at interactive rates have been demonstrated [193, 37], the significantly higher complexity of near-field reflectance fields currently requires too many images for every pair of viewing and illumination directions.

5

Acquiring Human Face Appearance

This section discusses the unique challenge of appearance acquisition of human faces and presents an example project that has leveraged principles of appearance acquisition and representation to acquire and model faces based on first principles, in order to analyze skin reflectance across many subjects. Creating digital faces that are indistinguishable from real ones is one of the biggest challenges in computer graphics. Although general rendering quality in graphics often achieves photo-realistic appearance of synthetic objects, rendering of human faces still remains a demanding task. This is not only because of the complexity of facial appearance, but also due to the fact that human observers are experts in judging whether a face looks “real” or not. The process of capturing an actor’s performance and likeness has aptly been named *digital face cloning* [11, 81]. Digital face cloning has many applications in movies, games, medicine, cosmetics, computer vision, biometrics, and virtual reality.

5.1 Skin Appearance

Particular attention has to be paid to the modeling of skin reflectance, as skin contributes the largest part of facial appearance. The dominant

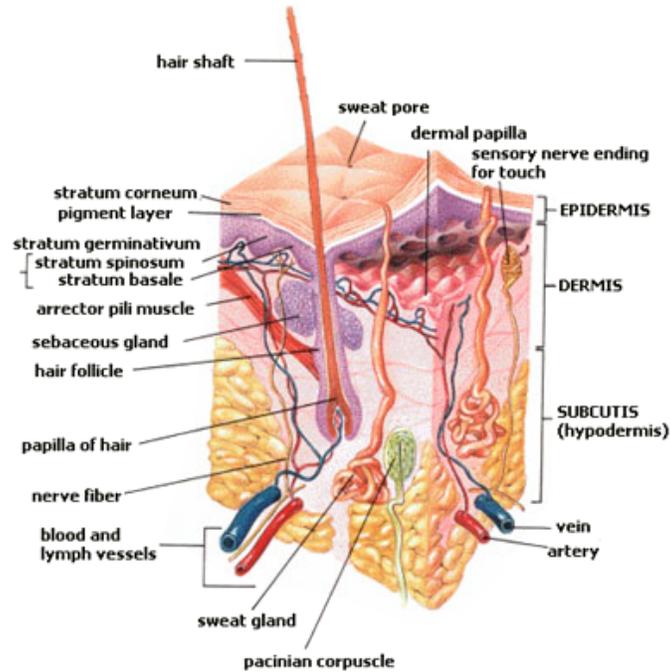


Fig. 5.1 Micro-scale, human skin is a very heterogeneous tissue. However, at scales relevant for rendering, it is sufficient to consider the two prominent layers: *epidermis* and *dermis*. The visual impact of scattering within the *hypodermis* is negligible. (Image from Wikipedia.)

effect in skin reflectance is due to skin's semi-translucent layers, see Figure 5.1. In a rough approximation, skin consists of two optically active layers, the epidermis and the dermis underneath it. Light transport is mainly affected by *surface reflection* at the air/skin interface (which can be described by a Bidirectional Reflectance-Distribution Function (BRDF)) and by internal scattering and absorption effects that take place after refraction at this interface. These internal effects comprise highly directional *single scattering* (light that exits the medium after a single, or very few scattering events) and mostly isotropic *multiple scattering* within and between the epidermal and dermal layers. Absorption is mainly affected by the distribution of *chromophores* in the skin, substances with characteristic absorption spectra whose distribution varies depending on anatomical and physiological conditions. An excellent survey on the physiological and anatomical

properties of skin and the state-of-the-art in skin appearance modeling has been published by Igarashi et al. [78].

Existing work on skin modeling employs all major classes of reflectance models, BRDF [65, 120, 170], Bidirectional Scattering-Surface Reflectance-Distribution Function (BSSRDF) [34, 35, 36, 83], and Bidirectional Texture Function (BTF) [21], see Section 2, to model skin reflectance. In practical environments, most skin models reduce skin reflectance to a two-component model, combining surface reflection and single scattering in a joint BRDF, and covering all internal scattering and absorption effects by a single diffusion approximation [11, 75, 160, 198]. More sophisticated, multi-layered, models exist [36, 55, 90], but they are either impossible or very demanding to be acquired for full faces. To date, the physiologically most advanced skin reflectance model that is still practical for rendering is arguably the model by Donner et al. [36]. They present efficient means to simulate heterogeneous subsurface scattering in a dual-layered skin model that is parameterized solely by chromophore distribution maps and demonstrate that the anatomically motivated model is capable of creating realistic skin appearance from simple, hand-drawn input parameter maps (Figure 5.2). They also present a multi-spectral acquisition procedure to obtain the model parameters from in-vivo measurements of skin, but the requirement of surface application of a water-based gel makes their acquisition method prohibitive for human faces.

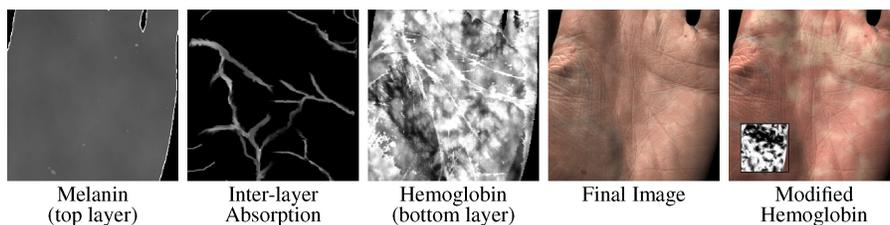


Fig. 5.2 Realistic skin appearance from hand-drawn parameter maps. Using a small set of simple, hand-drawn chromophore distribution maps and high-quality geometry of a palm (scan courtesy by Ma et al. [112]), the skin model by Donner et al. [36] generates realistic appearance and allows, for instance, to simulate appearance changes due to mechanical deformation: the right-most example shows a reduction of blood flow (hemoglobin concentration) after clenching and releasing the hand. The altered parameter map for hemoglobin is shown in the insets.

5.2 Face Acquisition Systems

Fueled by the steady interest in realistic face appearance in games and film industry, a variety of systems has been developed that aim at the appearance reproduction of a given subject's face.

Within the design space between parametric and non-parametric models (see Section 3.2), such systems traditionally have been closer to non-parametric, image-based modeling. This is motivated by the fact that image-based models implicitly capture effects such as self-shadowing, interreflections, and subsurface scattering, without the need to explicitly model every facial detail and its reflectance properties; they make it easier to achieve photo-realistic results and are less sensitive to measurement errors than explicit models. More recently, the general focus shifted toward describing face appearance from first principles by explicitly modeling facial geometry and surface properties, as, in general, explicit models are more directly accessible to editing operations, allow for more flexibility in post-production, and have a smaller memory footprint.

Pighin et al. [154] use a manually created impostor geometry and view-dependent texture mapping [32] to reproduce faces under static illumination conditions. Debevec et al. [30] use a *light stage* to acquire the dense reflectance field of a human face and present a process for creating realistic, relightable 3D face models by mapping image-based reflectance characteristics onto 3D-scanned geometry. While this method does consider the aggregate behavior of subsurface scattering in skin, it cannot reproduce correct subsurface scattering effects for heterogeneous illumination. Hawkins et al. [70] extend this technology to allow for deformations of the underlying face geometry, enabling changes of expression for the relightable face dataset. Wenger et al. [193] use rapidly time-multiplexed illumination to capture a reflectance field over time, allowing for relightable real-time performance, albeit only for a fixed viewpoint.

Borshukov and Lewis [11] combine an image-based model, high-quality explicit geometry, an analytic surface BRDF, and an image-space approximation for subsurface scattering to create highly realistic face models for a feature film. Fuchs et al. [46] use a lower-resolution

light stage to acquire reflectance fields from sparse viewing directions. While the reflectance field resolution would not be high enough to synthesize smooth lighting variations, they use clustering to determine component BRDFs (see Section 3.2.1). Weyrich et al. [195] use a much denser sampling (150 light sources, 16 cameras) to obtain enough reflectance samples to fit a BRDF in every surface point on the face. A separate contact-probe is used to acquire subsurface scattering properties of the skin. More detail on this project is given in Section 5.4. A contact-less approach to measuring subsurface scattering in a human face has been introduced by Tariq et al. [176], who use structured-light projection to densely sample the impulse response of the skin.

The latest peak in reconstruction quality of individual faces, and a largely parametric modeling approach, is the project presented by Ghosh et al. [55]. Using an elaborate combination of structured-light geometry acquisition, gradient illumination to determine high-quality normals [112], cross-polarization to identify surface reflectance ([30], see Section 2.5.1), active lighting to measure subsurface scattering [176], and global/local separation ([137], see Section 2.5.2) to separate localized scattering in the skin's upper regions from more de-localized scattering in deeper regions, their method acquires parameters for the multi-layered reflectance model shown in Figure 5.3. While

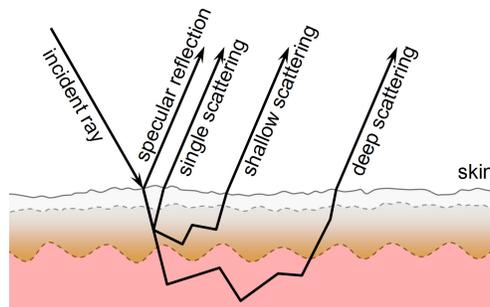


Fig. 5.3 Approximate skin reflectance components of the model by Ghosh et al. [55]. The individual components roughly correspond to the skin's layered structure and are assumed to be independent and purely additive. Surface reflectance is modeled by a microfacet-based BRDF model, single scattering using the first-order scattering BRDF model by Hanrahan and Krueger [65], and both subsurface scattering terms by the single-layered multi-pole BSSRDF model by Donner and Jensen [34]. (Image courtesy by Abhijeet Ghosh.)

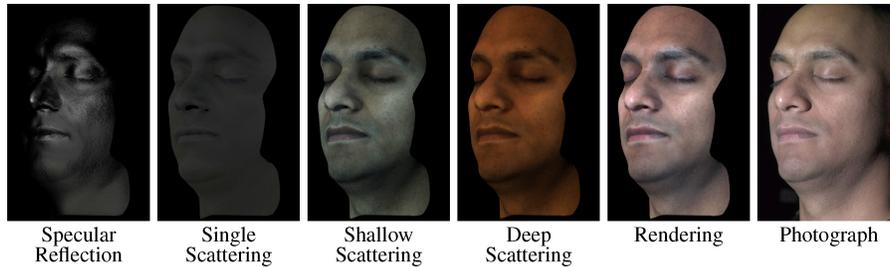


Fig. 5.4 Multi-layered face reconstruction by Ghosh et al. [55]. (Images courtesy by Abhijeet Ghosh.)

physiologically motivated, its components are purely additive, which may make the reflectance model physically less plausible. On the other hand, this design decision allows for an efficient fit to the measured data and leads to compelling reconstructions, as shown in Figure 5.4. The results also qualitatively capture the comparatively bluish contribution by the epidermis and the reddish color of the dermis, which is dominated by hemoglobin.

5.3 Appearance Editing

For any digital face representation, the ability to post-process the dataset and to edit facial appearance is desirable. Nevertheless, high-level appearance editing of general face datasets is still an underdeveloped field. Film productions, amongst the most prominent applications of digital face cloning, typically focus on face acquisition for specific shots, whose requirements are known at acquisition time, and use a significant amount of manual labor to meet the artistic requirements [11, 39, 75, 160, 198]. In particular, editing of facial appearance is still mostly within the realm of the artist, who manually edits facial textures. More automatic approaches to realistically alter appearance apply image-based techniques to photographs [107, 182]. A higher-level editing technique is to use morphable face models to blend between shape and texture of multiple individual face models [8, 46].

Generally, the goal of high-level appearance editing should be that the results continue to appear realistic to a human observer. It is to be expected that representations closer to the anatomical reality [36, 55]

bear the potential of guaranteeing realism under modifications. Even a realistic model, however, requires knowledge about realistic parameter ranges. The remainder of this chapter discusses a project that targets the analysis of facial reflectance to determine guidelines and high-level editing operations to alter facial appearance in a realistic way. An additional degree of detail is provided to illustrate how this project builds upon the principles discussed in the previous papers.

5.4 Analysis of Human Face Appearance

The face analysis project conducted by Weyrich et al. [194, 195] targeted the creation of a practical, explicit (non-image-based) representation for human faces, the creation of respective acquisition hardware, and the collection of reflectance data over a large population, in order to derive high-level parameters for realistic appearance editing.

5.4.1 Skin Reflectance Model

At its core is a simple skin reflectance model that aims to strike a balance between physical realism and practicability. Key requirements of such a model are generality, to allow for a flexible use, and editability, that is, it should be possible to change the face's appearance using intuitive controls. These design goals leave a wide range of potential models that are more or less strictly aligned with physical reality [78]. A key decision in the project, however, has been that the model should have a minimal number of degrees of freedom, to facilitate fitting the model to measured data. On the other hand, physical quantities within the model should be wellexposed to further an elementary understanding of skin reflectance.

The project uses a single layer to describe the bulk effect of subsurface scattering in epidermis and dermis, modulated by an additional modulation texture [57], in order to achieve spatial variation in the diffuse term. On top of the epidermis, a thin oil layer is assumed. This decomposition separates reflectance into specular and diffuse reflectance and is simple enough to be accessible to modeling from measured data.

Diffuse subsurface scattering is modeled using the dipole-diffusion approximation [83], see Section 2.4. This model assumes isotropic scattering (which, as will be shown, holds for facial skin) and a homogeneous material. The specular term is described by the widely-used Torrance–Sparrow BRDF model, which proved superior over other reflectance models in modeling skin gloss at the oily skin/air interface [194].

5.4.2 Data Acquisition

Two custom reflectance acquisition devices were built to capture surface and subsurface reflectance independently. A contact device allows for sparse subsurface measurements across the face. Instead of measuring pure surface reflectance, the second device captures the full skin reflectance as a whole. Skin reflectance being a sum of subsurface and surface components, surface reflectance can then be derived by subtracting an estimated subsurface component from the overall skin reflectance measurements.

5.4.2.1 Subsurface Scattering Acquisition

The first device, internally referred to as the “BSSRDF Gun”, is a contact-measurement device to measure purely subcutaneous light transport. The quantity of interest is the wavelength-dependent mean free path ℓ , or skin’s *translucency* $1/\ell$, respectively (see Section 2.4). The device feeds light into the skin using an optical fiber. A bundle of optical fibers arranged around the feeding fiber collects light exiting the skin at different distances to the feeding fiber. Digitizing the radiant exitance using an HDR camera at the end of the fiber bundle allows to measure the characteristic radial fall-off, that is, the diffusion kernel due to skin’s subsurface scattering. ℓ is obtained by fitting the dipole approximation of the diffusion model BSSRDF [83] to the measured fall-off.

As we will see in Section 5.4.4, translucency only varies minimally across the face. Hence, only a few measurements of the BSSRDF Gun are required to obtain representative translucency values of a face. Note that an even contact between the fiber probe and the skin is required

to eliminate surface reflectance effects. This is ensured by gently evacuating the sensor using a suction pump.

5.4.2.2 Reflectance Field Acquisition

The second measurement device samples skin reflectance as a whole, that is, the sum of subsurface and surface effects. The device captures a 150×16 reflectance field of the face within a spherical dome equipped with sixteen 1300×1030 firewire cameras and 150 disk-shaped LED panels. A commercial single-shot 3D scanner acquires the facial geometry, which is required to associate pixels in the reflectance field with surface points on the face. The setup is completely synchronized, that is, the 150 light sources are sequentially triggered, while all cameras simultaneously acquire images at 12 frames per second. Each lighting condition is imaged under two different exposure times to increase the dynamic range of the measurements. Hence, a full acquisition takes about 25 seconds.

The resulting reflectance field has to be radiometrically corrected, according to the spatial location of each surface point, considering distance to the respective light source and differences between camera sensitivities and light source characteristics (see Section 2.2). Figure 5.5 shows a sample reflectance field before correction. In order to improve the accuracy of the geometric reconstruction, the reflectance field is used to determine photometric normals that in turn are used to refine the geometry (Section 2.2.5).

5.4.3 Model Fit

The model fit determines model parameters for each surface point on the face. In a first step, the *diffuse albedo* R_d in each point is estimated by color-space-based separation of specular and diffuse reflection, see Section 2.5. This separation technique is extended to trade the Lambertian model of diffuse reflectance for a diffuse term that considers the transmissive Fresnel terms $F_t(\eta, \cdot)$, cf. (1.21). The same diffuse term occurs in the BRDF approximation of the dipole-diffusion model [83], see Section 1.4. The subsurface scattering BSSRDF is parameterized to meet the average diffuse albedo $\overline{R_d}$, while maintaining the translucency



Fig. 5.5 Raw reflectance images of a subject acquired with all 16 cameras and 14 (of 150) lighting conditions. Each row shows images with different camera viewpoints and the same lighting condition. Each column displays images with different lighting conditions and the same viewpoint. The images are not yet color corrected, revealing differences between camera characteristics and light source colors.

obtained using the “BSSRDF Gun”. The modulation texture is set to scale the underlying BSSRDF to meet R_d in every surface point. Finally, the parameters of the Torrance–Sparrow surface BRDF are obtained by fitting this BRDF to the residual reflectance samples after subtraction of the diffuse reflectance in each point.

5.4.4 Results and Analysis

After the model fit, model parameters for each surface point on the face are known. These parameters are encoded in floating-point textures over a common uv -parameterization of the facial geometry. Using custom shaders that implement the skin model within a Monte-Carlo raytracer, this allows to render photo-realistic images under arbitrary illumination from arbitrary vantage points. In particular, as the reflectance field has been acquired within a fully calibrated system, it becomes possible to replay the exact illumination and viewing conditions of each reflectance field image. This allows to directly evaluate our skin model in a side-by-side comparison with photographs from within the acquisition dome. Figure 5.6 shows such a comparison. It compares single input reflectance images with synthetic images for different faces and different viewpoints. Note, that a slight blur in the fitted model reflects the fact that each surface point's parameters are fitted to 2,400 input images simultaneously, which in the presence of noise and measurement imprecisions makes it impossible to retrieve the exact input image from the model.

A central goal of the project has been to analyze skin reflectance properties over a large group of subjects in order to obtain general insight into the variability of skin appearance across individuals.

5.4.4.1 The Face Database

To this end, 149 subjects were scanned and classified by skin type, gender, age, and other traits. Each scan was manually classified into facial regions, such as forehead, nose, chin, and others. This allows for a statistical analysis of characteristic variations in skin reflectance for different populations and across different facial regions. The skin type is classified according to the Fitzpatrick system [41]. Table 5.1 explains the Fitzpatrick system and shows the distribution of our measurements. Table 5.2 shows our face region classification.

5.4.4.2 Findings

Translucency Variance: An initial experiment on variation of skin translucency validated the isotropic assumption of the model's

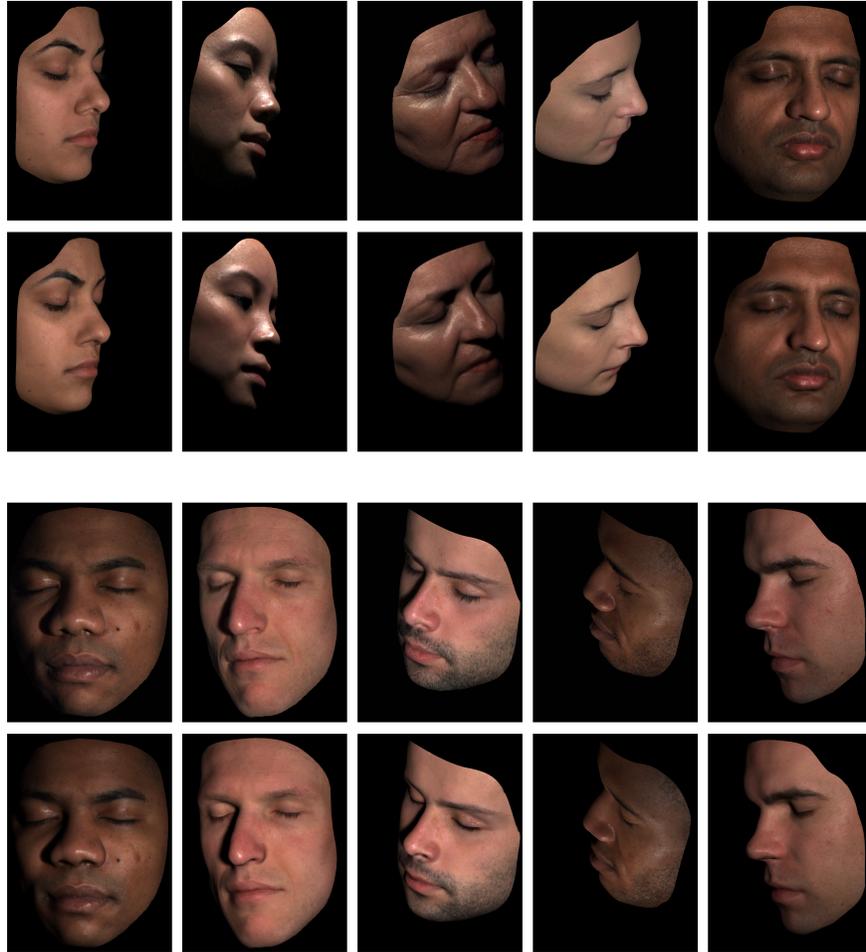


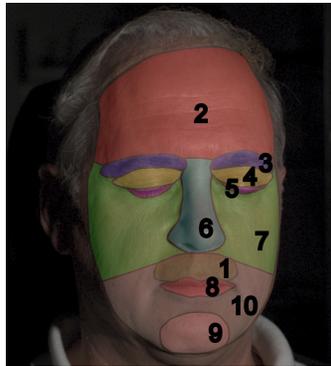
Fig. 5.6 Comparison of real photographs (first and third rows) to our model (second and last rows). All photographs were cropped according to the 3D model to remove distracting features.

subsurface scattering term. By taking subsurface scattering measurements under 16 different orientations of the sensor, the degree of anisotropy of the diffusion kernel has been measured. It turns out that light diffusion is not always isotropic; abdominal skin, for instance, shows a well-expressed scattering anisotropy. All facial measurements, however, show near-isotropic diffusion kernels, which justifies an isotropic BSSRDF model for facial skin. Spatial translucency variance

Table 5.1 The Fitzpatrick skin type system and the number of subjects per skin type.

Skin Type	Skin Color	Sun Exposure Reaction	Subjects (M/F)
I	Very white	Always burn	—
II	White	Usually burn	8 / 6
III	White to olive	Sometimes burn	49 / 18
IV	Brown	Rarely burn	40 / 8
V	Dark brown	Very rarely burn	13 / 2
VI	Black	Never burn	4 / 1

Table 5.2 Ten face regions: (1) mustache, (2) forehead, (3) eyebrows, (4) eyelids, (5) eye-lashes, (6) nose, (7) cheek, (8) lips, (9) chin, and (10) beard region.



was also analyzed by measuring 52 points in two subjects' faces. As far as accessible by the sensor, all facial regions showed a very uniform translucency. Ultimately, this led to the decision to model skin translucency to be constant across each face and reduce the number of measurements per subject to three. Analyzing translucency variations across multiple subjects revealed a subtle difference between male and female subjects (females having a slightly more translucent skin), while other traits did not correlate statistically significantly with translucency.

Spatial BRDF Variance: A more significant variability, however, could be found in the surface reflectance. The respective Torrance–Sparrow parameters ρ_s and m vary significantly in dependence of the facial region. For each facial region, principal component analysis (PCA) of these parameters has been performed, considering the BRDF fits of all subjects in the database. It turns out that the parameters do not only

vary between facial regions, but depending on the region, there is also a higher variability across subjects. Exemplary observations are: the nose is quite specular, while the chin is rather non-specular; the BRDF variance on the forehead is extremely low and almost uniform across subjects, while reflectance above the lip varies highly between subjects. This shows clearly that spatial BRDF variance is an important aspect of facial appearance.

Skin Trait Variance: In order to detect correlations between reflectance parameters and the traits associated with each subject, canonical correlation analysis (CCA) was performed. It turns out that the surface BRDF parameters correlate the most with skin type and gender. Less surprising, albedo is highly correlated with skin type. Apart from that, there is no significant correlation of albedo with other traits.

5.4.5 Appearance Transfer

It is now possible to use the parameter observations within the face database to derive intuitive user controls to alter facial appearance. While the analysis performed in the previous section can generally be used as a guideline when changing skin parameters, it is desirable to have higher-level controls. Weyrich et al. suggest the texture synthesis procedure by Heeger and Bergen [73] as a generic tool to transfer appearance parameters between subjects and to seamlessly blend between them [124]. The texture synthesis is applicable to all model parameter types and can be used to add freckles, moles, gloss variations, and other individual effects. With the face database at hand, this provides a general appearance editing framework. Figure 5.7 shows examples where this method has been applied to a face's diffuse reflectance, where changes are most visible. Altering other model parameters works analogously, although the effect appears to be more subtle in renderings.

5.4.6 Conclusion

This section presented a project that developed a simple and practical skin model. An important feature of this model is that all its parameters

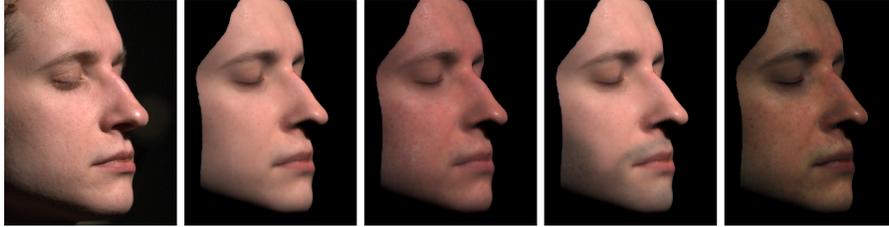


Fig. 5.7 Appearance editing, altering the diffuse reflectance. From left to right: Real photograph; rendering; making the face sun-burnt; adding hair follicles in the beard area; making the skin type darker.

can be robustly estimated from measurements. This reduces the large amount of measured data to a manageable size, facilitates editing, and enables face appearance changes. Images from the model come close to reproducing photographs of real faces for arbitrary illumination and pose. The model has been fit to data of a large and diverse group of people. The analysis of this data provides insight into the variance of face reflectance parameters based on age, gender, or skin type. The database with all statistics is available to the research community for face synthesis and analysis [130].

In general, the model itself may be extended to be closer to human anatomy, much in the lines of Ghosh et al. [55] or Donner et al. [36]. Despite its simplicity, however, it was already able to provide useful insight into skin reflectance in general. Other important areas that require a different modeling approach are facial hair (eyebrows, eyelashes, mustaches, and beards), hair, ears, eyes, and teeth. Very fine facial hair also leads to asperity scattering and the important “velvet” look of skin near grazing angles [88]. None of the models presented in this section take this into account.

6

Open Problems and Future Work

In this article we presented an introduction to appearance capture and suitable representations, together with the description of selected work that measures the spatially- and directionally-varying reflectance and subsurface scattering properties of complex materials. The presented topics have been selected to provide a good coverage of core problems in appearance acquisition. As any active research area, however, the field keeps on evolving, and there are various open problems currently under investigation. The remainder picks out a few topics that are representative for current research interests.

Practical Hardware Designs: Depending on the application scenario, existing hardware designs for appearance capture are not always practical. One challenge, for instance, is the design of *portable* measurement systems. An interesting approach has been explored by Kuthirummal and Nayar [92], who show that a radially-symmetric mirror placed in front of a camera can be used for, among other things, obtaining simultaneous measurements of multiple viewpoints for BRDF measurement. Ben-Ezra et al. [4] built a solid-state BRDF measurement device for rapid BRDF acquisition that consists of a dense set of LEDs in a hemispherical arrangement. The design exploits the fact that LEDs can be

used both as emitters and receivers, which enables a compact, portable system design and fast acquisition.

Interaction between wavelength and geometry: As mentioned in Section 1, it has been common to factor BRDF (and higher dimensional reflectance functions) into separated functions of wavelength and geometry:

$$f_r(\lambda, \omega_i \rightarrow \omega_o) \approx g(\lambda)f(\omega_i \rightarrow \omega_o). \quad (6.1)$$

It is not yet clear, however, to what degree this approximation is accurate. Indeed, as stated by Nicodemus et al. ([142], p. 31):

“Unfortunately, we are not aware of any data that will establish the extent to which there may be interaction between geometrical dependence and spectral dependence, except for the knowledge that it is a significant factor in some internal-scattering situations, as pointed out earlier, in addition to the obvious case of a diffraction grating where there is clearly substantial interaction. Certainly, there will be some wavelength bands over which the interaction is so slight that this last procedure will yield a good approximation, and it will probably hold more widely. But until more complete bidirectional reflectance-distribution data are gathered to establish such dependencies, we can only speculate and caution against the possibility of interaction effects.”

This situation has not changed significantly in the 30 years since this was written, and one interesting direction to pursue is the collection of multi-spectral reflection measurements that allow this interaction to be thoroughly explored.

Simultaneous recovery of shape and reflectance: Joint estimation of shape and reflectance is a challenge, compared to the traditional way of reflectance-independent estimation of geometry followed by reflectometry. Most examples of the former rely on parametric models [50, 59, 139, 174, 204], but with the right representations of reflectance this might be avoidable. In a practical system, the two approaches might be combined, with the latter used to provide initial conditions to bootstrap a joint estimation.

Acquisition ‘in the wild’: Many environments, especially those outdoors, do not permit precise lighting control. Estimating reflectance

(and shape) in these environments requires techniques that are different from many of those that have been discussed in this article. There are a number of approaches (both active and passive) that have been proposed to acquire appearance models in semi-controlled or uncontrolled environments, and almost all of them have relied on low-parameter BRDF models. Developing techniques that can handle a broader class of materials remain an open problem. Results in this direction seem imminent with the maturation of multiple-view geometry and camera autocalibration. There is also evidence that alternative non-parametric representations of BRDF exist that are more appropriate for this task [158]. It is also possible that there are additional physical constraints that have yet to be exploited [173].

The separation of reflection components can be especially useful in natural environments when controlled, active illumination is difficult or impossible. Local and non-local reflections for non-conducting surfaces can be separated (at least approximately) using a simple stick occluder under direct sunlight [137], and diffuse/specular separation can be achieved using color information according to the dichromatic model and/or passive polarization imaging [201, 138]. In many cases, passive polarization imaging can also be used for material identification, for example, to distinguish metallic from non-metallic surfaces [200, 201].

References

- [1] E. H. Adelson and J. R. Bergen, ch. The Plenoptic Function and Elements of Early Vision, pp. 3–20, *Computational Models of Visual Processing*. MIT Press, 1991.
- [2] N. Alldrin, T. Zickler, and D. Kriegman, “Photometric stereo with non-parametric and spatially-varying reflectance,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [3] M. Ashikhmin, S. Premoze, and P. Shirley, “A Microfacet-based BRDF Generator,” in *Computer Graphics (Proceedings of ACM SIGGRAPH 2001)*, pp. 65–74, 2001.
- [4] M. Ben-Ezra, J. Wang, B. Wilburn, X. Li, and L. Ma, “An LED-only BRDF measurement device,” in *Proceedings of IEEE CVPR 2008*, pp. 1–8, June 2008. IEEE Computer Society.
- [5] F. Bernardini and H. Rushmeier, “Strategies for registering range images from unknown camera positions,” *Proceedings of SPIE*, vol. 3958, pp. 200–206, 2000.
- [6] F. Bernardini and H. Rushmeier, “The 3D Model Acquisition Pipeline,” *Computer Graphics Forum*, vol. 21, no. 2, pp. 149–172, 2002.
- [7] F. Blais, “Review of 20 years of range sensor development,” *Journal of Electronic Imaging*, vol. 13, p. 231, 2004.
- [8] V. Blanz and T. Vetter, “A morphable model for the synthesis of 3D faces,” in *SIGGRAPH '99: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 187–194, New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1999.
- [9] J. F. Blinn, “Models of light reflection for computer synthesized pictures,” in *SIGGRAPH '77: Proceedings of the 4th Annual Conference on Computer*

- Graphics and Interactive Techniques*, pp. 192–198, New York, NY, USA: ACM Press, 1977.
- [10] S. Boivin and A. Gagalowicz, “Image-based rendering of diffuse, specular and glossy surfaces from a single image,” *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 107–116, 2001.
 - [11] G. Borshukov and J. P. Lewis, “Realistic human face rendering for “The Matrix Reloaded”,” in *ACM SIGGRAPH 2003 Conference Abstracts and Applications (Sketch)*, p. 1, New York, NY, USA: ACM Press, 2003.
 - [12] Camera Calibration Toolbox for Matlab. http://www.vision.caltech.edu/bouguetj/calib_doc/.
 - [13] S. Chandrasekhar, *Radiative Transfer*. Courier Dover Publications, 1960.
 - [14] B. Chen and H. P. A. Lensch, “Light source interpolation for sparsely sampled reflectance fields,” in *Vision, Modeling, and Visualization 2005 (VMV’05)*, (G. Greiner, J. Hornegger, H. Niemann, and M. Stamminger, eds.), pp. 461–469, Erlangen, Germany: Aka, November 2005.
 - [15] T. Chen, H. Lensch, C. Fuchs, and H. P. Seidel, “Polarization and phase-shifting for 3d scanning of translucent objects,” *Computer Vision and Pattern Recognition 2007. CVPR’07. IEEE Conference on*, pp. 1–8, 2007.
 - [16] H. Chong, S. Gortler, and T. Zickler, “The von Kries hypothesis and a basis for color constancy,” in *Proceedings of the International Conference on Computer Vision*, 2007.
 - [17] R. L. Cook, “Shade trees,” in *Computer Graphics SIGGRAPH 84 Proceedings*, pp. 223–231, 1984.
 - [18] R. L. Cook and K. E. Torrance, “A reflection model for computer graphics,” *ACM Transactions On Graphics*, vol. 1, no. 1, pp. 7–24, January 1982.
 - [19] O. Cula, K. Dana, F. Murphy, and B. Rao, “Bidirectional imaging and modeling of skin texture,” *IEEE Transactions on Biomedical Engineering*, vol. 51, no. 12, pp. 2148–2159, December 2004.
 - [20] O. G. Cula and K. Dana, “Compact representation of bidirectional texture functions,” in *Proceedings of Computer Vision and Pattern Recognition*, Kauai, HI, pages in press, December 2001.
 - [21] O. G. Cula and K. J. Dana, “Image-based Skin Analysis,” in *Texture 2002, The Second International Workshop on Texture Analysis and Synthesis*, pp. 35–42, Copenhagen, Denmark, June 2002.
 - [22] K. Dana, “BRDF/BTF Measurement device,” *ICCV*, pp. 460–466, 2001.
 - [23] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink, “Reflectance and texture of real-world surfaces,” in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 151–157, 1997.
 - [24] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink, “Reflectance and texture of real world surfaces,” *ACM Transactions on Graphics*, vol. 1, no. 18, pp. 1–34, 1999.
 - [25] K. J. Dana and J. Wang, “Device for convenient measurement of spatially varying bidirectional reflectance,” *Journal of the Optical Society of America A*, vol. 21, no. 1, pp. 1–12, 2004.
 - [26] K. Daubert, H. P. A. Lensch, W. Heidrich, and H.-P. Seidel, “Efficient cloth modeling and rendering,” in *Eurographics Rendering Workshop 2001*, 2001.

- [27] J. Davis and X. Chen, "Calibrating pan-tilt cameras in wide-area surveillance networks," in *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, pp. 144–149, 2003.
- [28] J. Davis, D. Nehab, R. Ramamoorthi, and S. Rusinkiewicz, "Spacetime stereo: A unifying framework for depth from triangulation," *Ieee Transactions on Pattern Analysis And Machine Intelligence*, pp. 296–302, 2005.
- [29] J. Davis, R. Ramamoorthi, and S. Rusinkiewicz, "Spacetime stereo: A unifying framework for depth from triangulation," *Proceedings of 2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003*, vol. 2, 2003.
- [30] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar, "Acquiring the reflectance field of a human face," in *Computer Graphics, SIGGRAPH 2000 Proceedings*, pp. 145–156, July 2000.
- [31] P. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *Computer Graphics, SIGGRAPH 97 Proceedings*, pp. 369–378, Los Angeles, CA, 1997.
- [32] P. Debevec, C. J. Taylor, and J. Malik, "Modeling and rendering architecture from photographs: A hybrid geometry- and image-based approach," in *Computer Graphics, SIGGRAPH 96 Proceedings*, pp. 11–20, August 1996.
- [33] P. Debevec, A. Wenger, C. Tchou, A. Gardner, J. Waese, and T. Hawkins, "A lighting reproduction approach to live-action compositing," *ACM Transactions on Graphics (SIGGRAPH 2002)*, vol. 21, no. 3, pp. 547–556, July 2002.
- [34] C. Donner and H. W. Jensen, "Light diffusion in multi-layered translucent materials," *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2005)*, vol. 24, no. 3, pp. 1032–1039, 2005.
- [35] C. Donner and H. W. Jensen, "A Spectral BSSRDF for shading human skin," in *Proceedings of the 16th Eurographics Symposium on Rendering*, pp. 409–417, Eurographics Association, June 2006.
- [36] C. Donner, T. Weyrich, E. d'Eon, R. Ramamoorthi, and S. Rusinkiewicz, "A layered, heterogeneous reflectance model for acquiring and rendering human skin," *To appear in ACM Transactions on Computer Graphics (Proceedings of SIGGRAPH ASIA)*, vol. 27, 2008.
- [37] P. Einarsson, C.-F. Chabert, A. Jones, W.-C. Ma, B. Lamond, T. Hawkins, M. Bolas, S. Sylwan, and P. Debevec, "Relighting human locomotion with flowed reflectance fields," in *Rendering Techniques 2006: 17th Eurographics Workshop on Rendering*, pp. 183–194, June 2006.
- [38] M. Fairchild, *Color Appearance Models*. Wiley-IS&T, 2nd Edition, 2005.
- [39] C. Feeny, "Servers, Spydercams and 'Spider-Man 2'," in *VFX Pro*, (http://www.uemedia.net/CPC/vfxpro/printer_9050.shtml), July 2004.
- [40] G. Finlayson, M. S. Drew, and B. V. Funt, "Color constancy: Generalized diagonal transforms suffice," *Journal of the Optical Society of America A: Optics and Image Science, and Vision*, vol. 11, no. 11, pp. 3011–3019, 1994.
- [41] T. B. Fitzpatrick, "The validity and practicality of sun-reactive skin types I through VI," *Archives of Dermatology*, vol. 124, pp. 869–871, 1988.

- [42] T. Franken, M. Dellepiane, F. Ganovelli, P. Cignoni, C. Montani, and R. Scopigno, "Minimizing user intervention in registering 2D images to 3D models," *The Visual Computer*, vol. 21, no. 8–10, pp. 619–628, 2005.
- [43] C. Fuchs, M. Goesele, T. Chen, and H.-P. Seidel, "An Empirical Model for Heterogeneous Translucent Objects," Research Report MPI-I-2005-4-006, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, May 2005.
- [44] M. Fuchs, V. Blanz, H. P. Lensch, and H.-P. Seidel, "Adaptive sampling of reflectance fields," *ACM Transactions on Graphics*, vol. 26, no. 2, pp. 10:1–10:18, June 2007.
- [45] M. Fuchs, H. P. A. Lensch, V. Blanz, and H.-P. Seidel, "Superresolution reflectance fields: Synthesizing images for intermediate light directions," *Computer Graphics Forum*, vol. 26, no. 3, pp. 447–456, September 2007.
- [46] M. Fuchs, H. P. A. Lensch, and H.-P. Seidel, "Reflectance from images: A model-based approach for human faces," *IEEE Transactions on Visualization and Computer Graphics*, vol. 11, no. 3, pp. 296–305, 2005.
- [47] R. Furukawa, H. Kawasaki, K. Ikeuchi, and M. Sakauchi, "Appearance based object modeling using texture database: Acquisition, compression and rendering," in *Proceedings of the 13th Eurographics Workshop on Rendering*, Pisa, Italy, June 2002.
- [48] A. Gardner, C. Tchou, T. Hawkins, and P. Debevec, "Linear light source reflectometry," *ACM Transactions on Graphics (SIGGRAPH 2003)*, vol. 22, no. 3, pp. 749–758, 2003.
- [49] G. Garg, E.-V. Talvala, M. Levoy, and H. P. A. Lensch, "Symmetric photography: Exploiting data-sparseness in reflectance fields," in *Rendering Techniques 2006: 17th Eurographics Workshop on Rendering*, pp. 251–262, June 2006.
- [50] A. Georgiades, "Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo," in *Proceedings of the Ninth IEEE International Conference on Computer Vision-Volume 2*, Washington, DC, USA: IEEE Computer Society, 2003.
- [51] A. Georgiades, P. Belhumeur, and D. Kriegman, "From few to many: Illumination Cone models for face recognition under variable lighting and pose," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 6, pp. 643–660, 2001.
- [52] A. S. Georgiades, "Recovering 3-D shape and reflectance from a small number of photographs," in *Proceedings of the 14th Eurographics workshop on Rendering*, pp. 230–240, Aire-la-Ville, Switzerland, Switzerland: Eurographics Association, 2003.
- [53] A. Gershun, "The Light Field, translated by P. Moon and G. Timoshenko," *Journal of Mathematical Physics*, vol. 18, no. 51, 1939.
- [54] A. Ghosh, S. Achutha, W. Heidrich, and M. O'Toole, "BRDF acquisition with basis illumination," *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pp. 1–8, 2007.
- [55] A. Ghosh, T. Hawkins, P. Peers, S. Frederiksen, and P. Debevec, "Practical modeling and acquisition of layered facial reflectance," *To appear in ACM Transactions on Computer Graphics (Proceedings of SIGGRAPH ASIA)*, vol. 27, 2008.

- [56] M. Goesele, H. Lensch, and H. Seidel, "Validation of color managed 3D appearance acquisition," in *Proceedings of 12th Color Imaging Conference (CIC12)*, 2004.
- [57] M. Goesele, H. P. A. Lensch, J. Lang, C. Fuchs, and H.-P. Seidel, "DISCO — Acquisition of translucent objects," *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2004)*, vol. 23, no. 3, 2004.
- [58] D. B. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz, "Shape and spatially-varying BRDFs from photometric stereo," Technical Report 04-05-03, University of Washington, 2003.
- [59] D. B. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz, "Shape and spatially-varying BRDFs from photometric stereo," in *IEEE International Conference on Computer Vision*, 2005.
- [60] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen, "The lumigraph," in *Computer Graphics, SIGGRAPH 96 Proceedings*, pp. 43–54, New Orleans, LS, August 1996.
- [61] M. Grossberg and S. K. Nayar, "Modeling the space of camera response functions," *IEEE Transactions on Pattern Analysis And Machine Intelligence*, pp. 1272–1282, 2004.
- [62] J. Gu, C.-I. Tu, R. Ramamoorthi, P. Belhumeur, W. Matusik, and S. Nayar, "Time-varying surface appearance: Acquisition, modeling and rendering," *ACM Transactions on Graphics (SIGGRAPH 2006)*, vol. 25, no. 3, 2006.
- [63] W. Hackbusch, "A sparse matrix arithmetic based on \mathcal{H} -matrices. Part I: Introduction to \mathcal{H} -matrices," *Computing*, vol. 62, no. 2, pp. 89–108, 1999.
- [64] J. Y. Han and K. Perlin, "Measuring bidirectional texture reflectance with a kaleidoscope," *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 741–748, July 2003.
- [65] P. Hanrahan and W. Krueger, "Reflection from layered surfaces due to subsurface scattering," in *Computer Graphics, SIGGRAPH 93 Proceedings*, pp. 165–174, Anaheim, CA, August 1993.
- [66] B. Hapke, "A theoretical photometric function for the lunar surface," *Journal of Geophysical Research*, vol. 68, no. 15, 1963.
- [67] K. Hara, K. Nishino, and K. Ikeuchi, "Mixture of spherical distributions for single-view relighting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 1, pp. 25–35, 2008.
- [68] J. A. Hartigan and M. A. Wong, "A K-means clustering algorithm," *Applied Statistics*, vol. 28, pp. 100–108, 1979.
- [69] T. Hawkins, P. Einarsson, and P. Debevec, "A dual light stage," in *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, pp. 91–98, June 2005.
- [70] T. Hawkins, A. Wenger, C. Tchou, A. Gardner, F. Göransson, and P. Debevec, "Animatable facial reflectance fields," in *Rendering Techniques '04 (Proceedings of the Second Eurographics Symposium on Rendering)*, pp. 309–320, Norrköping, Sweden, June 2004.
- [71] X. D. He, K. E. Torrance, F. X. Sillion, and D. P. Greenberg, "A comprehensive physical model for light reflection," *Computer Graphics*, vol. 25, no. Annual Conference Series, pp. 175–186, 1991.

- [72] G. Healey, “Using color for geometry-insensitive segmentation,” *Journal of Optical Society of America A*, vol. 6, no. 6, pp. 920–937, 1989.
- [73] D. J. Heeger and J. R. Bergen, “Pyramid-based texture analysis/synthesis,” in *Proceedings of SIGGRAPH 95, Computer Graphics Proceedings, Annual Conference Series*, pp. 229–238, August 1995.
- [74] L. C. Henyey and J. L. Greenstein, “Diffuse radiation in the Galaxy,” *The Astrophysical Journal*, vol. 93, p. 70, 1941.
- [75] C. Héry, “Face cloning at ILM,” in *SIGGRAPH 2003 Course “Digital Face Cloning”*, 2003.
- [76] T. Hofmann, “Probabilistic latent semantic analysis,” in *Proceedings of Uncertainty in Artificial Intelligence*, 1999.
- [77] M. Holroyd, J. Lawrence, G. Humphreys, and T. Zickler, “A photometric approach for estimating normals and tangents,” *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 2008.
- [78] T. Igarashi, K. Nishino, and S. K. Nayar, “The appearance of human skin: A survey,” *Foundations and Trends[®] in Computer Graphics and Vision*, vol. 3, no. 1, pp. 1–95, 2007.
- [79] I. Ihrke, K. N. Kutulakos, H. P. A. Lensch, M. Magnor, and W. Heidrich, “State of the art in transparent and specular object reconstruction,” in *STAR Proceedings of Eurographics*, 2008.
- [80] A. Ishimaru, *Electromagnetic Wave Propagation, Radiation, and Scattering*. Prentice Hall, 1991.
- [81] H. W. Jensen, “Digital face cloning,” in *SIGGRAPH ’03: ACM SIGGRAPH 2003 Sketches & Applications*, pp. 1–1, New York, NY, USA: ACM, 2003.
- [82] H. W. Jensen and J. Buhler, “A rapid hierarchical rendering technique for translucent materials,” in *Computer Graphics, SIGGRAPH 2002 Proceedings*, pp. 576–581, Los Angeles, CA, July 2002.
- [83] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan, “A practical model for subsurface light transport,” in *Computer Graphics, SIGGRAPH 2001 Proceedings*, pp. 511–518, Los Angeles, CA, August 2001.
- [84] A. Jones, A. Gardner, M. Bolas, I. McDowall, and P. Debevec, “Performance geometry capture for spatially varying relighting,” *ACM SIGGRAPH 2005 Conference Abstracts and Applications (Sketch)*, 2005.
- [85] K. Karner, H. Mayer, and M. Gervautz, “An image based measurement system for anisotropic reflection,” *Computer Graphics Forum*, vol. 15, no. 3, pp. 119–128, 1996.
- [86] J. Kautz and M. D. McCool, “Interactive rendering with arbitrary BRDFs using separable approximations,” in *Rendering Techniques ’99 (Proceedings of the Tenth Eurographics Workshop on Rendering)*, pp. 281–292, New York, NY: Springer Wien, June 1999.
- [87] G. Klinker, S. Shafer, and T. Kanade, “A physical approach to color image understanding,” *International Journal of Computer Vision*, vol. 4, no. 1, pp. 7–38, 1990.
- [88] J. Koenderink and S. Pont, “The secret of velvety skin,” *Machine Vision and Application*, vol. 14, pp. 260–268, Special Issue on Human Modeling, Analysis, and Synthesis, 2003.

- [89] M. L. Koudelka, S. Magda, P. N. Belhumeur, and D. J. Kriegman, "Acquisition, compression, and synthesis of bidirectional texture functions," in *Proceedings of the 3rd International Workshop on Texture Analysis and Synthesis*, 2003.
- [90] A. Krishnaswamy and G. V. G. Baranoski, "A biophysically-based spectral model of light interaction with human skin," *Computer Graphics Forum*, vol. 23, no. 3, pp. 331–340, September 2004.
- [91] R. Kurazume, K. Nishino, M. D. Wheeler, and K. Ikeuchi, "Mapping textures on 3D geometric model using reflectance image," *Systems and Computers in Japan*, vol. 36, no. 13, pp. 92–101, 2005.
- [92] S. Kuthirummal and S. K. Nayar, "Multiview radial catadioptric imaging for scene capture," *Proceedings of ACM SIGGRAPH 2006*, vol. 25, no. 3, pp. 916–923, 2006.
- [93] E. P. F. Lafortune, S.-C. Foo, K. E. Torrance, and D. P. Greenberg, "Non-linear approximation of reflectance functions," in *SIGGRAPH '97: Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 117–126, New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1997.
- [94] J. Lawrence, A. Ben-Artzi, C. DeCoro, W. Matusik, H. Pfister, R. Ramamoorthi, and S. Rusinkiewicz, "Inverse shade trees for non-parametric material representation and editing," *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2006)*, vol. 25, no. 3, 2006.
- [95] J. Lawrence, S. Rusinkiewicz, and R. Ramamoorthi, "Efficient BRDF importance sampling using a factored representation," *ACM Transactions on Graphics (SIGGRAPH 2004)*, vol. 23, no. 3, pp. 496–505, August 2004.
- [96] D. Lee and H. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, pp. 788–791, 1999.
- [97] H. C. Lee, E. J. Breneman, and C. P. Schulte, "Modeling light reflection for computer color vision," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 12, no. 4, pp. 402–409, 1990.
- [98] H. P. A. Lensch, W. Heidrich, and H.-P. Seidel, "Automated texture registration and stitching for real world surfaces," in *Proceedings of the Pacific Graphics Conference*, 2000.
- [99] H. P. A. Lensch, W. Heidrich, and H.-P. Seidel, "A silhouette-based algorithm for texture registration and stitching," *Graphical Models*, vol. 64, no. 4, pp. 245–262, 2001.
- [100] H. P. A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel, "Image-based reconstruction of spatially varying materials," in *Proceedings of the 12th Eurographics Workshop on Rendering*, pp. 104–115, June 2001.
- [101] H. P. A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel, "Image-based reconstruction of spatial appearance and geometric detail," *ACM Transactions on Graphics*, vol. 22, no. 2, pp. 234–257, 2003.
- [102] T. Leung and J. Malik, "Representing and recognizing the visual appearance of materials using three-dimensional textons," *International Journal of Computer Vision*, vol. 43, no. 1, pp. 29–44, 2001.

- [103] K. Levenberg, "A method for the solution of certain non-linear problems in least squares," *The Quarterly of Applied Mathematics*, vol. 2, pp. 164–168, 1944.
- [104] M. Levoy, "The Stanford spherical gantry," <http://graphics.stanford.edu/projects/gantry/>. Last accessed August 16, 2008.
- [105] M. Levoy and P. Hanrahan, "Light field rendering," in *Computer Graphics, SIGGRAPH 96 Proceedings*, pp. 31–42, New Orleans, LS, August 1996.
- [106] R. R. Lewis, "Making shaders more physically plausible," *Computer Graphics Forum*, vol. 13, no. 2, 1994.
- [107] T. Leyvand, D. Cohen-Or, G. Dror, and D. Lischinski, "Data-driven enhancement of facial attractiveness," *ACM Transactions on Graphics (SIGGRAPH 2008)*, vol. 27, no. 3, pp. 1–9, 2008.
- [108] H. Li, S. C. Foo, K. E. Torrance, and S. H. Westin, "Automated three-axis gonioreflectometer for computer graphics applications," *Optical Engineering*, vol. 45, p. 043605, 2006.
- [109] X. Liu, Y. Hu, J. Zhang, X. Tong, B. Guo, and H.-Y. Shum, "Synthesis and rendering of bidirectional texture functions on arbitrary surfaces," *IEEE Transactions on Visualization and Computer Graphics*, vol. 10, no. 3, pp. 278–289, 2004.
- [110] J. Lu, A. S. Georghiadis, A. Glaser, H. Wu, L.-Y. Wei, B. Guo, J. Dorsey, and H. Rushmeier, "Context-aware textures," *ACM Transactions on Graphics*, vol. 26, no. 1, p. 3, 2007.
- [111] R. Lu, J. Koenderink, and A. Kappers, "Optical Properties (Bidirectional Reflectance Distribution Functions) of Velvet," *Applied Optics*, vol. 37, no. 25, pp. 5974–5984, September 1998.
- [112] W.-C. Ma, T. Hawkins, P. Peers, C.-F. Chabert, M. Weiss, and P. Debevec, "Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination," *Eurographics Symposium on Rendering*, pp. 183–194, 2007.
- [113] S. Magda and D. Kriegman, "Reconstruction of volumetric surface textures for real-time rendering," in *Proceedings of the Eurographics Symposium on Rendering*, 2006.
- [114] S. P. Mallick, T. Zickler, D. J. Kriegman, and P. N. Belhumeur, "Specularity removal in images and videos: A PDE approach," in *European Conference on Computer Vision*, pp. 550–563, 2006.
- [115] S. P. Mallick, T. E. Zickler, D. J. Kriegman, and P. N. Belhumeur, "Beyond lambert: Reconstructing specular surfaces using color," in *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) — Volume 2*, pp. 619–626, Washington, DC, USA: IEEE Computer Society, 2005.
- [116] T. Malzbender, D. Gelb, and H. Wolters, "Polynomial texture maps," in *Computer Graphics, SIGGRAPH 2001 Proceedings*, pp. 519–528, Los Angeles, CA, 2001.
- [117] S. R. Marschner, "Inverse rendering for computer graphics," PhD thesis, Cornell University, Ithaca, NY, 1998.

- [118] S. R. Marschner, A. Arbrece, , and J. T. Moon, “Measuring and modeling the appearance of wood,” *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2005)*, vol. 24, no. 3, 2005.
- [119] S. R. Marschner, S. H. Westin, E. P. F. Lafortune, and K. E. Torrance, “Image-based measurement of the Bidirectional Reflectance Distribution Function,” *Applied Optics*, vol. 39, no. 16, pp. 2592–2600, June 2000.
- [120] S. R. Marschner, S. H. Westin, E. P. F. Lafortune, K. E. Torrance, and D. P. Greenberg, “Image-based BRDF measurement including human skin,” in *Proceedings of the 10th Eurographics Workshop on Rendering*, pp. 139–152, Granada, Spain, June 1999.
- [121] V. Masselus, P. Peers, P. Dutré, and Y. D. Willems, “Relighting with 4D incident light fields,” *ACM Transactions on Graphics*, vol. 22, no. 3, pp. 613–620, July 2003.
- [122] W. Matusik, H. Pfister, M. Brand, and L. McMillan, “A data-driven reflectance model,” *ACM Transactions on Graphics (SIGGRAPH 2003)*, vol. 22, no. 3, pp. 759–770, July 2003.
- [123] W. Matusik, H. Pfister, M. Brand, and L. McMillan, “Efficient isotropic BRDF measurement,” *Proceedings of the 14th Eurographics Workshop on Rendering*, pp. 241–247, 2003.
- [124] W. Matusik, M. Zwicker, and F. Durand, “Texture design using a simplicial complex of morphable textures,” *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 787–794, 2005.
- [125] D. McAllister, A. Lastra, and W. Heidrich, “Efficient rendering of spatial bi-directional reflectance distribution functions,” *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware*, pp. 79–88, 2002.
- [126] D. K. McAllister, “A generalized surface appearance representation for computer graphics,” PhD thesis, University of North Carolina (UNC), Chapel Hill, NC, 2002.
- [127] M. D. McCool, J. Ang, and A. Ahmad, “Homomorphic factorization of BRDFs for high-performance rendering,” *Computer Graphics*, vol. 35, no. Annual Conference Series, pp. 171–178, 2001.
- [128] S. Mersch, “Polarized lighting for machine vision applications,” in *Proceedings of RI/SME Third Annual Applied Machine Vision Conference*, pp. 40–54, Schaumburg, IL, February 1984.
- [129] J. Meseth, G. Müller, M. Sattler, R. Sarlette, A. Artusi, A. Wilkie, G. Zotti, R. Klein, and W. Purgathofer, “High quality rendering of reflectance data,” CGI 2004 tutorial, June 2004.
- [130] Mitsubishi Electric Research Laboratories (MERL), ETH Zurich. The MERL/ETH Skin Reflectance Database. <http://www.merl.com/facescanning/>.
- [131] T. Mitsunaga and S. Nayar, “Radiometric self calibration,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, pp. 374–380, June 1999.
- [132] G. Müller, G. H. Bendels, and R. Klein, “Rapid synchronous acquisition of geometry and appearance of cultural heritage artefacts,” in *VAST 2005: 6th*

- International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*, pp. 13–20, November 2005.
- [133] G. Muller, J. Meseth, and R. Klein, “Compression and real-time rendering of measured BTFs using local PCA,” *Vision, Modeling and Visualization*, 2003.
- [134] G. Müller, J. Meseth, M. Sattler, R. Sarlette, and R. Klein, “Acquisition, synthesis and rendering of bidirectional texture functions,” *Computer Graphics Forum*, vol. 24, no. 1, 2005.
- [135] G. Muller, J. Meseth, M. Sattler, R. Sarlette, and R. Klein, “Acquisition, synthesis and rendering of the bidirectional texture function,” *Computer Graphics Forum*, vol. 24, no. 1, 2005.
- [136] S. Nayar, K. Ikeuchi, and T. Kanade, “Surface reflection: physical and geometrical perspectives,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 611–634, 1991.
- [137] S. Nayar, G. Krishnan, M. Grossberg, and R. Raskar, “Fast separation of direct and global components of a scene using high frequency illumination,” *ACM Transactions on Graphics (TOG)*, vol. 25, no. 3, pp. 935–944, 2006.
- [138] S. K. Nayar, X. Fang, and T. Boult, “Separation of reflection components using color and polarization,” *International Journal of Computer Vision*, vol. 21, no. 3, pp. 163–186, 1997.
- [139] S. K. Nayar, K. Ikeuchi, and T. Kanade, “Determining shape and reflectance of hybrid surfaces by photometric sampling,” *Robotics and Automation, IEEE Transactions on*, vol. 6, no. 4, pp. 418–431, 1990.
- [140] D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi, “Efficiently combining positions and normals for precise 3D geometry,” *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2005)*, vol. 24, no. 3, pp. 536–543, 2005.
- [141] A. Ngan, F. Durand, and W. Matusik, “Experimental analysis of BRDF models,” in *Proceedings of the 15th Eurographics Symposium on Rendering*, pp. 117–226, Eurographics Association, 2005.
- [142] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis, “Geometric considerations and nomenclature for reflectance,” Monograph 161, National Bureau of Standards (US), October 1977.
- [143] K. Nishino, Y. Sato, and K. Ikeuchi, “Eigen-texture method: Appearance compression based on 3D model,” in *Proceedings of Computer Vision and Pattern Recognition*, pp. 618–624, June 1999.
- [144] K. Nishino, Z. Zhang, and K. Ikeuchi, “Determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image synthesis,” in *Proceedings of IEEE International Conference on Computer Vision ICCV*, pp. 599–606, July 2001.
- [145] Open Source Computer Vision Library. <http://opencvlibrary.sourceforge.net/>.
- [146] M. Oren and S. Nayar, “Generalization of Lambert’s Reflectance Model,” in *Proceedings of SIGGRAPH*, 1994.
- [147] G. Patow and X. Pueyo, “A survey of inverse rendering problems,” *Computer Graphics Forum*, vol. 22, no. 4, pp. 663–687, 2003.
- [148] F. Pedrotti and L. Pedrotti, *Introduction to Optics*. Benjamin Cummings, 2nd Edition, 1993.

- [149] P. Peers, N. Tamura, W. Matusik, and P. Debevec, "Post-production facial performance relighting using reflectance transfer," *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, vol. 26, no. 3, pp. 52:1–52:10, July 2007.
- [150] P. Peers, K. vom Berge, W. Matusik, R. Ramamoorthi, J. Lawrence, S. Rusinkiewicz, and P. Dutré, "A compact factored representation of heterogeneous subsurface scattering," *ACM Transactions on Graphics*, vol. 25, no. 3, pp. 746–753, July 2006.
- [151] F. Pellacini and J. Lawrence, "AppWand: Editing measured materials using appearance-driven optimization," *ACM Transactions on Graphics (SIGGRAPH 2007)*, vol. 26, no. 3, 2007.
- [152] M. Pharr and P. Hanrahan, "Monte Carlo evaluation of non-linear scattering equations for subsurface reflection," in *Computer Graphics, SIGGRAPH 2000 Proceedings*, pp. 75–84, New Orleans, LS, July 2000.
- [153] B. T. Phong, "Illumination for computer generated pictures," *Communications of the ACM*, vol. 18, no. 6, pp. 311–317, June 1975.
- [154] F. Pighin, J. Hecker, D. Lischinski, R. Szeliski, and D. H. Salesin, "Synthesizing realistic facial expressions from photographs," in *Computer Graphics*, vol. 32 of *SIGGRAPH 98 Proceedings*, pp. 75–84, 1998.
- [155] R. Ramamoorthi and P. Hanrahan, "A signal-processing framework for inverse rendering," in *Computer Graphics, SIGGRAPH 2001 Proceedings*, pp. 117–128, 2001.
- [156] R. Raskar, M. Brown, R. Yang, W. Chen, G. Welch, H. Towles, B. Seales, and H. Fuchs, "Multi-projector displays using camera-based registration," in *Proceedings of IEEE Conference on Visualization*, pp. 161–168, 1999.
- [157] C. Rocchini, P. Cignoni, C. Montani, and R. Scopigno, "Acquiring, stitching and blending diffuse appearance attributes on 3D models," *The Visual Computer*, vol. 18, no. 3, pp. 186–204, 2002.
- [158] F. Romeiro, Y. Vasilyev, and T. Zickler, "Passive reflectometry," in *ECCV '08: Proceedings of the 10th European Conference on Computer Vision*, pp. 859–872, London, UK: Springer-Verlag, 2008.
- [159] S. Rusinkiewicz, "A new change of variables for efficient BRDF representation," in *Rendering Techniques '98 (Proceedings of Eurographics Rendering Workshop '98)*, (G. Drettakis and N. Max, eds.), pp. 11–22, New York, NY: Springer Wien, 1998.
- [160] M. Sagar, "Reflectance field rendering of human faces for "Spider-Man 2"," in *ACM SIGGRAPH 2004 Sketches*, August 2004.
- [161] Y. Sato and K. Ikeuchi, "Temporal-color space analysis of reflection," *Journal of the Optical Society of America A*, vol. 11, no. 11, pp. 2990–3002, 1994.
- [162] Y. Sato, M. D. Wheeler, and K. Ikeuchi, "Object shape and reflectance modeling from observation," in *Computer Graphics, SIGGRAPH 97 Proceedings*, pp. 379–387, 1997.
- [163] M. Sattler, R. Sarlette, and R. Klein, "Efficient and Realistic Visualization of Cloth," in *Eurographics Symposium on Rendering 2003*, June 2003.
- [164] Y. Schechner, S. K. Nayar, and P. N. Belhumeur, "Multiplexing for optimal lighting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1339–1354, 2007.

- [165] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski, "A comparison and evaluation of multi-view stereo reconstruction algorithms," *International Conference on Computer Vision and Pattern Recognition*, pp. 519–528, 2006.
- [166] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch, "Dual photography," *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 745–755, August 2005.
- [167] S. Shafer, "Using color to separate reflection components," *COLOR Research and Applications*, vol. 10, no. 4, pp. 210–218, 1985.
- [168] A. Shashua and T. Riklin-Raviv, "The quotient image: Class-based re-rendering and recognition with varying illuminations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 2, pp. 129–139, 2001.
- [169] Specification ICC.1:2004-10 (Profile version 4.2.0.0). Available at: <http://www.color.org>, 2004.
- [170] J. Stam, "An illumination model for a skin layer bounded by rough surfaces," in *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, pp. 39–52, Vienna, Austria, London, UK: Springer-Verlag, June 2001.
- [171] A. Standard, "E1392-90, Standard practice for angle resolved optical scatter measurements on specular or diffuse surfaces," *American Society for Testing and Materials*.
- [172] M. M. Stark, J. Arvo, and B. Smits, "Barycentric Parameterizations for Isotropic BRDFs," *IEEE Transactions on Visualization and Computer Graphics*, vol. 11, no. 2, pp. 126–138, 2005.
- [173] K. Sunkavalli, F. Romeiro, W. Matusik, T. Zickler, and H. Pfister, "What do color changes reveal about an outdoor scene?," *IEEE Conference on Computer Vision and Pattern Recognition, 2008. CVPR 2008*, pp. 1–8, June 2008.
- [174] H. D. Tagare and R. J. P. de Figueiredo, "A theory of photometric stereo for a class of diffuse non-Lambertian surfaces," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 2, pp. 133–152, 1991.
- [175] R. T. Tan and K. Ikeuchi, "Separating reflection components of textured surfaces using a single image," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 2, pp. 178–193, 2005.
- [176] S. Tariq, A. Gardner, I. Llamas, A. Jones, P. Debevec, and G. Turk, "Efficient estimation of spatially varying subsurface scattering parameters," Technical Report ICT-TR-01-2006 July 27, 2006, University of Southern California, July 2006.
- [177] J. Tenenbaum, V. de Silva, and J. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, no. 5500, pp. 2319–2323, December 2000.
- [178] S. Tominga and B. A. Wandell, "Standard surface-reflectance model and illuminant estimation," *Journal of the Optical Society of America A*, vol. 6, no. 4, pp. 576–584, April 1989.
- [179] X. Tong, J. Wang, S. Lin, B. Guo, and H.-Y. Shum, "Modeling and rendering of quasi-homogeneous materials," *ACM Transactions on Graphics (SIGGRAPH 2005)*, vol. 24, no. 3, pp. 1054–1061, 2005.

- [180] X. Tong, J. Zhang, L. Liu, Xi. Wang, B. Guo, and H.-Y. Shum, "Synthesis of bidirectional texture functions on arbitrary surfaces," in *Proceedings of SIGGRAPH*, pp. 665–672, ACM Press, 2002.
- [181] K. E. Torrance and E. M. Sparrow, "Theory for off-specular reflection from roughened surfaces," *JOSA*, vol. 57, no. 9, pp. 1105–1114, September 1967.
- [182] N. Tsumura, N. Ojima, K. Sato, M. Shiraishi, H. Shimizu, H. Nabeshima, S. Akazaki, K. Hori, and Y. Miyake, "Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin," *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2003)*, vol. 22, no. 3, pp. 770–779, 2003.
- [183] H. Unten and K. Ikeuchi, "Color alignment in texture mapping under point light source and general lighting condition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.
- [184] M. A. O. Vasilescu and D. Terzopoulos, "TensorTextures: Multilinear image-based rendering," *ACM Transactions on Graphics (SIGGRAPH 2004)*, vol. 24, no. 3, 2004.
- [185] P. Viola and W. M. Wells, "Alignment by Maximization of Mutual Information," *International Journal of Computer Vision*, vol. 24, no. 2, pp. 137–154, 1997.
- [186] H. von Helmholtz, *Handbuch der Physiologischen Optik*. Leipzig: Leopold Voss, 1867.
- [187] H. Wang, Q. Wu, L. Shi, Y. Yu, and N. Ahuja, "Out-of-core tensor approximation of multi-dimensional matrices of visual data," *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2005)*, vol. 24, no. 3, pp. 527–535, 2005.
- [188] J. Wang, X. Tong, S. Lin, H. Bao, B. Guo, and H.-Y. Shum, "Appearance manifolds for modeling time-variant appearance of materials," *ACM Transactions on Graphics (SIGGRAPH 2006)*, vol. 25, no. 3, 2006.
- [189] J. Wang, S. Zhao, X. Tong, S. Lin, Z. Lin, Y. Dong, B. Guo, and H.-Y. Shum, "Modeling and rendering of heterogeneous translucent materials using the diffusion equation," *ACM Transactions on Graphics*, vol. 27, no. 1, pp. 1–18, 2008.
- [190] J. Wang, S. Zhao, X. Tong, J. Snyder, and B. Guo, "Modeling anisotropic surface reflectance with example-based microfacet synthesis," *ACM Transactions on Graphics (SIGGRAPH 2008)*, vol. 27, no. 3, pp. 41:1–41:9, 2008.
- [191] G. Ward, "Measuring and modeling anisotropic reflection," *Computer Graphics*, vol. 26, no. Annual Conference Series, pp. 265–273, 1992.
- [192] R. P. Weistroffer, K. R. Walcott, G. Humphreys, and J. Lawrence, "Efficient basis decomposition for scattered reflectance data," in *Proceedings of the Eurographics Symposium on Rendering*, Eurographics Association, 2007.
- [193] A. Wenger, A. Gardner, C. Tchou, J. Unger, T. Hawkins, and P. Debevec, "Performance relighting and reflectance transformation with time-multiplexed illumination," *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 756–764, August 2005.
- [194] T. Weyrich, "Acquisition of human faces using a measurement-based skin reflectance model," PhD thesis, Department of Computer Science, ETH Zurich, 2006.

- [195] T. Weyrich, W. Matusik, H. Pfister, B. Bickel, C. Donner, C. Tu, J. McAndless, J. Lee, A. Ngan, H. W. Jensen, and M. Gross, "Analysis of Human Faces using a Measurement-Based Skin Reflectance Model," *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2006)*, pp. 1013–1024, July 2006.
- [196] T. Weyrich, W. Matusik, H. Pfister, J. Lee, A. Ngan, H. W. Jensen, and M. Gross, "A measurement-based skin reflectance model for face rendering and editing," Technical Report, Mitsubishi Electric Research Laboratories, TR2005-071, December 2005.
- [197] D. R. White, P. Saunders, S. J. Bonsey, J. van de Ven, and H. Edgar, "Reflectometer for measuring the bidirectional reflectance of rough surfaces," *Applied Optics*, vol. 37, pp. 3450–3454, 1998.
- [198] L. Williams, "Case study: The gemini man," in *SIGGRAPH 2005 Course 'Digital Face Cloning'*, 2005.
- [199] N. Williams, K.-L. Low, C. Hantak, M. Pollefeys, and A. Lastra, "Automatic image alignment for 3D environment modeling," in *Proceedings of the Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI)*, 2004.
- [200] L. Wolff, "Polarization Methods in Computer Vision," PhD thesis, Columbia University, 1991.
- [201] L. B. Wolff and T. E. Boult, "Constraining object features using a polarization reflectance model," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 635–657, 1991.
- [202] D. N. Wood, D. I. Azuma, K. Aldinger, B. Curless, T. Duchamp, D. H. Salesin, and W. Stuetzle, "Surface light fields for 3D photography," in *Computer Graphics, SIGGRAPH 2000 Proceedings*, pp. 287–296, Los Angeles, CA, July 2000.
- [203] R. Woodham, "Photometric stereo: A reflectance map technique for determining surface orientation from image intensity," in *Proceedings of SPIE*, vol. 153, pp. 136–143, 1978.
- [204] K.-J. Yoon, A. Delaunoy, P. Gargallo, and P. Sturm, "Toward global and model based multiview stereo methods for shape and reflectance estimation," in *Proceedings of the First International Workshop on Photometric Analysis For Computer Vision (in conjunction with ICCV'07)*, 2007.
- [205] T. Yu, N. Xu, and N. Ahuja, "Recovering shape and reflectance model of non-lambertian objects from multiple views," *Proceedings of Computer Vision and Pattern Recognition (CVPR)*, 2004.
- [206] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse global illumination: Recovering reflectance models of real scenes from photographs," in *Computer Graphics, SIGGRAPH 99 Proceedings*, pp. 215–224, Los Angeles, CA, August 1999.
- [207] A. Zalesny and L. Van Gool, "Multiview texture models," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001.
- [208] L. Zhang, B. Curless, and S. Seitz, "Spacetime stereo: Shape recovery for dynamic scenes," *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, 2003.

- [209] L. Zhang and S. Nayar, “Projection defocus analysis for scene capture and image display,” *Proceedings of ACM SIGGRAPH 2006*, vol. 25, no. 3, pp. 907–915, 2006.
- [210] Z. Zhang, “Flexible camera calibration by viewing a plane from unknown orientations,” *International Conference on Computer Vision*, vol. 1, pp. 666–673, 1999.
- [211] T. Zickler, P. N. Belhumeur, and D. J. Kriegman, “Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction,” in *ECCV '02: Proceedings of the 7th European Conference on Computer Vision-Part III*, pp. 869–884, London, UK: Springer-Verlag, 2002.
- [212] T. Zickler, R. Ramamoorthi, S. Enrique, and P. Belhumeur, “Reflectance sharing: Predicting appearance from a sparse set of images of a known shape,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 8, pp. 1287–1302, 2006.